

Top 10 Windows 8.1 Apps

Windows IT Pro

A PENTON PUBLICATION

MARCH 2014 | WINDOWSITPRO.COM | WE'RE IN IT WITH YOU

Microsoft's **Perry Clarke** on the Future of Exchange Server

Explore Group-Object
in PowerShell

File Classification Infrastructure
in Windows Server 2012

System Center 2012
Orchestrator

Windows Azure Backup

Need to Know:
Windows 8.1 Update 1
and Windows Phone 8.1



STORAGECRAFT®

GRANULAR RECOVERY FOR EXCHANGE

Email Search and Rescue

Easily search, recover and migrate email.

Managing Exchange servers can be a challenge because of complexity. But when email servers are unavailable—even for a few minutes—productivity suffers. Ensure your email is always recoverable with StorageCraft® Granular Recovery for Exchange.

Simplify Complex Exchange Tasks: StorageCraft Granular Recovery for Exchange simplifies complex Exchange database management tasks with its simple drag-and-drop interface so you can:

- Search, recover and migrate emails, mailboxes or other mailbox items
- Complete basic and advanced searches
- Consolidate Exchange servers using the exclusive migration feature—including consolidation of Exchange servers running different versions of Exchange
- Enjoy fast enumeration of Exchange data—historically a very time consuming task

Fully Compatible with Microsoft Exchange: StorageCraft Granular Recovery for Exchange is fully compatible with the recent Microsoft® Exchange 2013 as well as Microsoft Exchange 2003, 2007, and 2010. Plus, you can use it with any backup solution including StorageCraft backup images.

Try StorageCraft Granular Recovery for Exchange for yourself by downloading a FREE 30-day trial today!



Download your
FREE 30-Day Trial
www.StorageCraft.com/WINGRE




STORAGECRAFT®
Backup Fast, Recover Faster




When turning it off and on again isn't an option.

Ontrack PowerControls 7.1


Search.Collect.Restore. **Made Simple.**



800.645.3649 | www.ontrackpowercontrols.com

 **Kroll Ontrack.**

COVER STORY



29

Talking Exchange Server with Microsoft's Perry Clarke

— Tony Redmond

Perry Clarke, corporate VP for Microsoft Exchange, and Tony Redmond discuss the issues surrounding Exchange Server, including the cloud, the future of on-premises software, and how mobile computing has changed the IT landscape.

Features

44 Windows Server 2012 File Classification Infrastructure
John Savill

59 Exploring PowerShell's Group-Object Cmdlet
Jeffery Hicks

70 Getting Started with System Center 2012 Orchestrator
John Savill

Products

87 New & Improved

Interact

82 Ask the Experts

In Every Issue

91 Advertiser Directory

91 Directory of Services

91 Vendor Directory

Chat with Us



Facebook



Twitter



LinkedIn

Columns



9

[Need to Know](#)

Windows 8.1 Update 1 and Windows Phone 8.1

Paul Thurrott



15

[Windows Power Tools](#)

Going Further with ForEach

Mark Minasi



18

[Top 10](#)

Top 10 Windows 8.1 Apps

Michael Otey



21

[What Would Microsoft Support Do?](#)

Introducing Windows Azure Backup

Robert Mitchell

Editorial

Vice President, Content & User Engagement:
Joe Panettieri

Editorial Director: Megan Keller

Editor-in-Chief: Amy Eisenberg

Senior Technical Director: Michael Otey

Technical Director: Sean Deuby

Senior Technical Analyst: Paul Thurrott

IT Community Manager: Rod Trent

Systems Management, Exchange Server &
Outlook: Jason Bovberg

Scripting, Developer Content:

Blair Greenwood

SQL Server: Jayleen Heft

SharePoint, Active Directory, Security,

Virtualization: Caroline Marwitz

Managing Editor: Lavon Peters

Senior Contributing Editors

David Chernicoff, Mark Minasi,
Tony Redmond, Paul Robichaux,
Mark Russinovich, John Savill

Contributing Editors

Alex K. Angelopoulos, Michael Dragone,
Jeff Felling, Brett Hill, Dan Holme,
Darren Mar-Elia, Eric B. Rux,
William Sheldon, Curt Spanburgh,
Bill Stewart, Orin Thomas, Douglas Toombs,
Ethan Wilansky

Art & Production

Senior Graphic Designer: Matt Wiebe
Group Production Manager:
Julie Jantzer-Ward
Project Manager: Adriane Wineinger

Advertising Sales

Strategic Accounts Director:
Chrissy Ferraro • 970-203-2883

Account Executives:

Megan Key • 970-203-2844

Barbara Ritter • 858-367-8058

Cass Schulz • 858-357-7649

Client Services

Senior Client Services Manager:
Michelle Andrews • 970-613-4964

Marketing & Circulation

Customer Service • 800-793-5697

Vice President, User Marketing &

Marketing Analytics: Tricia Syed

Technology Division & Penton Marketing Services

Senior Vice President: Sanjay Mutha

Corporate

Chief Executive Officer:

David Kieselstein

Chief Financial Officer/Executive Vice
President: Nicola Allais



List Rentals

Sarah Nowowiejski

Reprints

Reprint Sales:

Wright's Media • 877-652-5295

Windows IT Pro, March 2014, Issue No. 235,
ISSN 1552-3136. *Windows IT Pro* is published monthly by
Penton. Copyright ©2014 Penton. All rights reserved. No
part of this publication may be reproduced or distributed
in any way without the written consent of Penton.

Windows IT Pro, 748 Whalers Way, Fort Collins, CO 80525,
800-621-1544 or 970-663-4700. Customer Service:
800-793-5697.

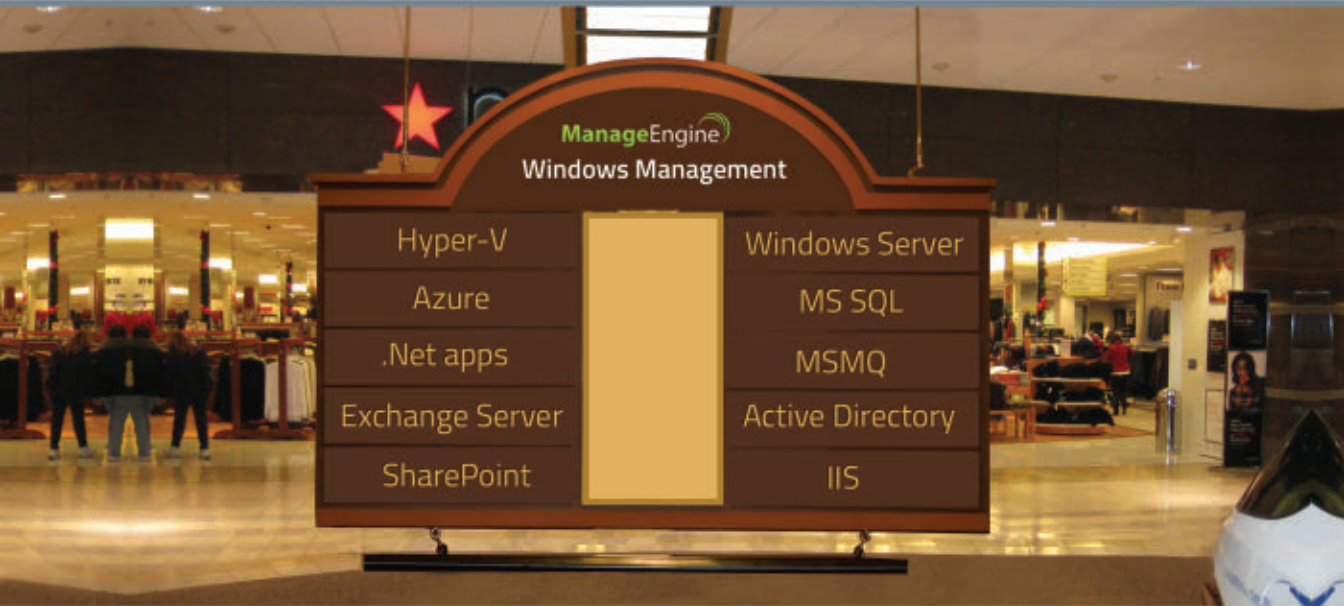
We welcome your comments and suggestions about the
content of *Windows IT Pro*. We reserve the right to edit all
submissions. Letters should include your name and
address. Please direct all letters to letters@windowsitpro.com. IT pros interested in writing for *Windows IT Pro* can
submit articles at [windowsitpro.com/node/submission/](http://windowsitpro.com/node/submission/article)
article.

Program Code: Unless otherwise noted, all programming
code in this issue is ©2014, Penton, all rights reserved.
These programs may not be reproduced or distributed
in any form without permission in writing from the
publisher. It is the reader's responsibility to ensure
procedures and techniques used from this publication are
accurate and appropriate for the user's installation. No
warranty is implied or expressed.

Windows®, Windows Vista®, and Windows Server®
are trademarks or registered trademarks of Microsoft
Corporation in the United States and/or other countries
and are used by Penton, under license from owner.
Windows IT Pro is an independent publication not
affiliated with Microsoft Corporation. Microsoft
Corporation is not responsible in any way for the editorial
policy or other contents of the publication.

Windows IT Pro

MANAGE AND OPTIMIZE YOUR WINDOWS APPS. **ALL FROM ONE PLACE.**



Gartner.
APM Magic Quadrant



Windows
Networking.com

Gold Award
★★★★

Why look around when you can get all your Windows management needs in one place? ManageEngine Applications Manager provides reliable Windows performance management with out-of-the-box monitors that are easy to set up and manage.

Applications Manager provides total visibility into the health, availability and performance of your Windows infrastructure at any moment of the day. Over 5000 customers agree.



+1 - 925 - 924 - 9500
www.manageengine.com/apm
demo.appmanager.com

ManageEngine
Applications Manager

Audit Like a Boss with Netwrix Auditor

Take Control of Change Auditing for



Active Directory



Exchange



File Servers



Windows Server



SQL Server



VMware

- On-demand reporting
- State-in-time reporting
- Compliance auditing

Start Your Free Trial:
netwrix.com/be_the_boss



PASS BUSINESS ANALYTICS CONFERENCE

May 7-9, 2014 | San Jose, CA



The World of Data is Changing. Stay ahead of the curve

Looking to improve your career path? Keep up with the rapidly changing analytics landscape and immerse yourself in 2 full days of technical sessions from top analytics experts, including:

intuit.

 Microsoft

 SurveyMonkey

 WELLS
FARGO

YAHOO!

"As a relative beginner to business analytics it was a bit like drinking from the fire hose but that's what I came for!"



Join acclaimed data visualization expert, author and TedTalk speaker **David McCandless** as he takes you on a journey through the world of visualizing facts, data, ideas, and statistics during the Keynote.

www.passbaconference.com

Sign up today and get \$200 off by using code **BAC200ITP** when you register.



Windows 8.1 Update 1 and Windows Phone 8.1

Information has recently emerged about Windows 9, code-named Threshold, which is due in April 2015. But that's still more than a year away, and Microsoft plans to deliver other Windows updates in the interim, updates which—despite their apparently minor version numbers—are arguably just as important for anyone using Windows.

Windows 8.1 Update 1

Microsoft plans to release its first major update for Windows 8.1 as soon as early March; this update adopts the servicing model and naming convention of Windows Phone. This makes sense, since Microsoft's newly minted Operating Systems business unit is run by Terry Myerson who—wait for it—used to run Windows Phone (before Steve Ballmer blew up the old corporate structure as part of a massive corporate reorganization).

Windows 8.1 Update 1, as this new update is called, can be seen as a step toward a consolidation of Microsoft's client OSs, a much-needed change that seems to be the overreaching goal of the Myerson era. As such, it might not seem far-reaching, but it does continue many of the themes seen in the initial Windows 8.1 release while providing a peek at the direction from which updates will continue.

With Windows 8.1, Microsoft sought to quiet its critics and appease those customers—in particular, corporate ones—who were none too pleased with the mobile-first philosophy of Windows 8—while not stepping back from the strategy of making Windows a first-class mobile experience. In this, they were successful: Windows 8.1 makes it possible for users on traditional PC hardware to mostly ignore the “Metro” mobile interfaces that bother them, while allowing those brave few who have adopted new Windows tablets, 2-in-1s, and



Paul Thurrott

is senior technical analyst for *Windows IT Pro*. He writes the SuperSite for Windows, a weekly editorial for *Windows IT Pro UPDATE*, and a daily Windows news and information newsletter called *WinInfo Daily UPDATE*.



Email



Twitter



Website

other hybrid PCs with multi-touch capabilities to keep their desktop interactions to a minimum.

If you really think about what happened in Windows 8.1, you might see the underlying problem: Sure, the system works better on new and old PC designs than did Windows 8, but it doesn't help the OS seem cohesive, with two operating environments—or “personalities”—that work together. The jarring gulf between the desktop and Metro is even wider in Windows 8.1 than it was in its predecessor.

We already know that [Windows 9 will help bridge the desktop and Metro worlds](#) by letting Metro mobile apps run on the desktop in windowed modes, and by returning—as an option—some form of a Start menu. But [Windows 9 won't ship to customers until April 2015, Microsoft expects](#), and its year-long gestation suggests that it won't be a major release, but rather something akin to Windows 8.1. My sources tell me the Windows 9 naming decision is based more around a desire to distance the OS from the poorly received Windows 8.

So if Windows 8.1—and, later, Windows 9—are evidence of what the OS team can accomplish in one year of development—and they are—we might look at Windows 8.1 Update 1 as what the team can accomplish in less than half that time. The total development time of Update 1 might be closer to four months than six months. (This, too, seems related to Microsoft's schedule for Windows Phone 8 updates: The firm shipped three, the last called Update 3, in about a year.)

What might those accomplishments be, specifically? Continuing within the theme of a more consolidated future, Update 1 will add the ability to pin Metro mobile apps to the desktop taskbar so that they can be launched from there. (You can currently pin both types of apps, Metro and desktop, to the Start screen.) But since windowed mobile apps won't become available until Windows 9, those pinned apps will still run full-screen, as they do now, and obscure both the desktop and its taskbar.

Microsoft is also adding right-click context menus to the Metro environment to make it friendlier to mouse-equipped users on traditional

PCs. It's not currently clear whether this capability is limited to the Start and Apps screens, or if it will work in Metro apps as well.

If you're familiar with these interfaces today, you know that right-clicking in Metro displays app bars that are more suitable for a multi-touch interface. So using context menus when a mouse is clicked will make the system more usable and familiar for those users. Responding to complaints about the new placement of power options (i.e., shutdown, sleep, and other settings) in Windows 8, Microsoft is also adding a prominent power button to the Start screen.

Microsoft is also utilizing a technology—which one unconfirmed source told me is called “wimboot”—to help decrease the amount of on-disk storage that's required by the OS. This is a big deal on today's modern Windows devices, which often ship with just 32GB or 64GB of SSD or similar storage, much of which is taken up by the OS installation, and, on many devices, by Office. But that same unconfirmed source told me that this disk savings will be offered only on new PCs and isn't something that can be added otherwise.

Windows 8.1 Update 1 will be finalized sometime in March and will be delivered via Windows Update, first as an optional update. I'm told that the broad public rollout will happen in early April.

Naturally, talk of an Update 1, combined with an understanding of how Microsoft updated Windows Phone 8 last year, leads to talk about whether the firm will deliver further updates—Update 2, for example—between Update 1 and Windows 9. That's possible, of course. But multiple sources have told me that the OS group—like much of Microsoft—is currently at something of a standstill while they wait for the reorganization to be completed (at the time this article was written, Satya Nadella had just been named CEO).

Also unclear is what Windows Phone 8.1 Update 1 does to further Microsoft's aim of consolidating Windows and Windows Phone. It's not hard to imagine Windows Phone and Windows RT merging in the future.

Are there any hints of this change coming in Update 1? It doesn't appear so. But Microsoft is also prepping a Windows Phone 8.1

release for the same basic time frame as Update 1—April 2014—and *that* release, allegedly, will begin that work.

Windows Phone 8.1

Much less is known about the next major update for Windows Phone, which is behind schedule despite an 18-month development window. In fact, Mr. Myerson pulled people off of Windows to help get Windows Phone 8.1 delivered in time for its April release. I’m told it’s going to be close and that 8.1 is “coming in hot.”

I’m told that Microsoft will announce 8.1 at the Build conference in April, but I’d be surprised if we didn’t see some Windows Phone 8.1 news earlier. We know that existing phones running Windows Phone 8 are upgradeable, though we can expect carriers to limit upgrades.

Beyond that, what we know about the expected Windows Phone 8.1 feature set is mostly rumor, at least from my perspective. There is a Siri/Google Now–like voice-control feature that’s code-named Cortana (and should just be called Cortana), after a character in the “Halo” games. Microsoft will let hardware makers remove the three front-facing buttons—Back, Start, and Search—and use new software-based versions, as we see on Android. And many of the integrated experiences, or hubs, should be redesigned.

We also know that a so-called [Enterprise Feature Pack](#) (expect a name change) is coming to Windows Phone early this year, though it’s not clear if this is tied to Windows Phone 8.1. This update will add features such as S/MIME email encryption support, an app-aware, auto-triggered VPN like the one in Windows 8.1, enterprise Wi-Fi support with EAP-TLS, certificate management, and enhanced Mobile Device Management (MDM) policies so you can lock down the phone in more granular ways than are possible with Exchange ActiveSync (EAS).

Why These Updates Matter

Microsoft’s update naming conventions have certainly changed over the years. The service packs and feature packs of years past gave way

to cumulative update rollups, general distribution releases (GDRs), and now these Update 1-type releases. So what? The servicing and updating of Windows has always been of interest and a challenge for those who need to manage those changes across multiple desktops, right? Well, yes. But two things are different now.

One, Microsoft would like to get its OSs—and other major products such as Office—to the point where they are essentially auto-updating all the time and doing so in a way that doesn't disrupt things even in highly managed environments. And two, for perhaps the first time ever, and certainly to this degree, forces external to Microsoft are driving this change.

Windows 8 itself was a reaction to the multi-touch and tablet-based personal computing world that Microsoft saw coming right around the time that Apple announced its first iPad. (I know, I know. You might argue that Windows 8—and Surface—is really just a *reaction* to the iPad, and although Microsoft claims otherwise, let's not quibble over the timing.)

And whether you love or loathe this new Windows, one point isn't debatable: Personal computing has absolutely changed in the ensuing years, with individuals turning to simpler and less complex devices. This trend will continue and accelerate.

As of the end of 2013, over 92 percent of all PCs sold ran Windows, and if that number sounds surprisingly familiar, well, you're right: Windows hasn't lost hold of this important market, ever. But PC sales have fallen. And in the emerging new market for personal computing, in which smartphones, tablets, PCs, and PC-like devices must all be supported and counted as equals, Windows is no longer dominant.

In 2013, hardware makers shipped over 1 billion smartphones to customers, far more than the roughly 300 million PCs that shipped in the same time frame. Eighty percent of those devices run Android, while 15 percent run Apple's iOS and only 4 percent or so utilize Windows Phone. Microsoft's smartphone OS is seeing some success—it garnered over 10 percent market share in Europe last year, for

example, and even outsells iPhone in some markets—but it's coming at the low end of the market and there are questions about how well the [Nokia](#) transition will go.

Tablet sales, meanwhile, are growing, and hardware makers delivered 220 million of the devices in 2013. Tablet sales will exceed those of PCs as soon as this year. Microsoft's share of this market, so far, is likewise nonexistent. By 2017, Windows is expected to command just 10 percent of the tablet market, compared to almost 60 percent for Android and 30 percent for Apple's iPad.

To manage the transition to the future, Microsoft will step up the development of mobile apps on competing platforms, but a key goal is to strengthen the performance of Windows in markets such as smartphones and tablets. This means moving away from the monolithic, three-year development time that it undertook for Windows versions in the past.

Windows and Windows Phone today are imperfect, and they are ill-equipped for the markets in which they compete. So releases such as Windows Phone 8.1 and Windows 8.1 Update 1—as well as subsequent updates that keep current OS versions fresh and updated with new features—represent one way in which Microsoft can try to compete more effectively with its mobile OS competitors. So is listening to customers and offering user experiences that work well on each of these very different device types.

The next couple of years are going to be very important for Windows. Although we might still get excited about the Big Bang releases, success now hinges on smaller, incremental updates. And on whether Microsoft's customers—especially the corporate base that makes up two thirds of its business—accept these kinds of changes. ■

Going Further with ForEach

Create better PowerShell one-liners

In “Introducing the Pipeline and ForEach” and “From One-Liner to ForEach One-Liner,” I’ve shown you how the ForEach cmdlet and the pipeline variable `$_` together let you write simple one-liners in a different way. For example, suppose bigfirm.com has an organizational unit (OU) called Machinists, and you want to change the *description* attribute of every user account to the text *Machinist*. (Yes, it’s a trivial example, but it’s also easy to understand.) You could do that with this simple one-liner, of the kind that I discussed a dozen or more columns ago in “Find Users with Get-ADUser”:

```
get-aduser -filter * -searchbase "ou=machinists,dc=bigfirm,dc=com" | set-aduser -description "Machinist"
```

And as I demonstrated in last month’s column, you can recast that one-liner into one using ForEach and `$_`, like so:

```
get-aduser -filter * -searchbase "ou=machinists,dc=bigfirm,dc=com" | foreach {set-aduser $_ -description "Machinist"}
```

That second approach accomplishes exactly what the first one does, but it’s more complicated. So why do it? Simple! It’s more flexible. Let’s say you want to set every user’s description not to *Machinist* but to something whimsical. You want to extract each user’s first name and set his or her description to *[First Name] the Machinist* (e.g., *Dan the Machinist*). In that case, the simpler, non-ForEach one-liners I’ve shown you so far would fail—there’s just no way to do it.

The first time I tried to do something like that, my reasoning was as follows:



Mark Minasi

is a senior contributing editor for *Windows IT Pro*, an MCSE, and the author of 30 books, including *Mastering Windows Server 2008 R2* (Sybex). He writes and speaks around the world about Windows networking.



Email



Twitter



Website

1. The Active Directory (AD) attribute containing a user's first name is called *givenname*.
2. It's an attribute, so I can retrieve it by doing some kind of `Get-ADUser` command and adding *.givenname*, as in this example, which retrieves the first name of a user named *jdorn*:

```
$someuser = (get-aduser jdorn)
$someuser.givenname
```

3. Things in the pipeline already live in a variable (the built-in variable `$_`), and after *get-aduser* there's a user in the pipeline, so *\$_givenname* should contain that user's first name.
4. As I've demonstrated in past columns, I can use the plus sign (+) to glue two text strings together (the programmer term is *concatenate*), so if I have a user account in the pipeline with a first name of Dan, PowerShell should be able to create *Dan the Machinist* with the text

```
$_givenname + "the Machinist"
```

5. So, this ought to work:

```
get-aduser -filter * -searchbase "ou=machinists,dc=bigfirm,dc=com" | set-aduser -description ($_givenname + "the Machinist")
```

Although my reasoning was sound, it didn't work in PowerShell. Why not? The one-liners in this column so far have been powerful because many AD-oriented PowerShell one-liners accomplish what used to require 50 lines of VBScript. PowerShell masks the complexity of decision-making and looping. But it can only take it so far. The basic rule is that whenever you want to use `$_` somewhere on the right side of the pipeline, you need to pull out `ForEach` and a scriptblock.

So, what would make it work? Simply re-cast the simple one-liner format into the mildly more complex ForEach-oriented one. Just take my looks-good-but-doesn't-work one-liner, rebuild it with ForEach and \$_, and you get the now-working one-liner

```
get-aduser -filter * -searchbase "ou=machinists,dc=bigfirm,
dc=com" | foreach {set-aduser $_ -description ($_.givenname
+ "the Machinist")}
```

Using ForEach doesn't have to be hard. Just remember that most simple one-liners look like *filter, then hammer*, as in

```
get-aduser | unlock-adaccount
```

To make that kind of construction ForEach-friendly, follow these steps:

1. Leave the filter (the part to the left of the pipeline) unchanged.
2. Leave the pipeline in place.
3. Type *foreach* { to start the scriptblock.
4. The hammer cmdlet on the right side of the pipeline almost certainly needs to know the name of the user account that you want it to modify, but in the ForEach world it doesn't automatically see the object in the pipeline. Therefore, present the user account name as \$_. What was simply *unlock-adaccount* must now become *unlock-adaccount \$_*.
5. Type the remainder of the command, if applicable, and now you're free to use \$_.whatever to extract properties from the user account in the pipeline, such as the *givenname* example in this column. (Oh, and recall that Get-ADUser returns only about 10 of an AD user's dozens of properties by default. If you want everything, add *-properties ** to get them all.)
6. Don't forget to end the script block with }.

Next month, you'll cook up a ForEach one-liner of your own! ■

Top 10 Windows 8.1 Apps

These apps can help you in the workplace



**Michael
Otey**

is senior technical director
for *Windows IT Pro* and
SQL Server Pro.

Email



When all is said and done, apps—or maybe the lack thereof—will stand out as the factor that makes or breaks [Windows 8.1/8](#). When you create a new interface, you need a rich ecosystem of apps to enrich the platform. Unfortunately, after more than a year since the release of Windows 8, Microsoft is still working to populate the Windows Store with Windows 8 apps. As of December 2013, there were 130,892 apps in the Windows Store. Although that's nothing to sneeze at, the number is dwarfed by the more than 1 million apps in the Apple Store and the more than 700,000 apps in Google Play. But if you look around, you'll find a number of cool apps that are more useful to IT pros than *Angry Birds*. These are my top 10 Windows 8.1 apps.

⑩ Skype

One of the most useful apps for Windows 8.1 is the Skype app. As you might expect, the Skype app lets you IM and send files, as well as make and receive video calls. In Windows 8.1, the Skype app provides the ability to answer calls from the lock screen without having to log on. The new [Skype app](#) is built in to Windows 8.1, but you can also get it for free from the Windows Store.

⑨ Facebook

It might be a stretch to include Facebook in a list of apps for IT pros, but you know you use it. The Facebook app lets you check and post status updates and photos, as well as chat with other Facebook users. The Share charm allows you to automatically attach the photos you're viewing to Facebook posts. You can download the [Facebook app](#) for free from the Windows Store.

8 Twitter

The official Twitter app's look and feel resemble that of the Twitter web page. The app supports URL shortening, multiple Twitter accounts, and picture attachments. The [Twitter app](#) is built in to Windows 8.1, but you can also download it for free from the Windows Store.

7 Tweetium

If you're a big Twitter user, you might want to check out Tweetium as an alternative to the Twitter app. Hailed by many as the best Twitter client, Tweetium is laid out a bit like TweetDeck. With Tweetium, your tweets, messages, and notifications are presented horizontally. You can block users or specific hashtags from your feed, and an indicator tells you when you last used the app. [Tweetium](#) is \$2.99 in the Windows Store.

6 Nextgen Reader

An RSS reader for Windows 8 and Windows Phone, Nextgen Reader allows you to pin multiple live tiles to the Start screen; it displays a list of feeds along the left side of the screen and a reading pane on the right. Tapping the app's logo in the top left corner shifts it into the Windows 8 tiled interface. [Nextgen Reader](#) is \$2.99 in the Windows Store, but a free trial is also available.

5 Flipboard

Flipboard is a personal magazine that you create by collecting search results for people, topics, hashtags, blogs, and other websites, which Flipboard organizes into a digital magazine format. The Flipboard app can search through Twitter, Facebook, Instagram, and Google + . Flipboard supports live tiles and looks great in Windows 8.1. You can get [Flipboard](#) for free from the Windows Store.

4 Adobe Photoshop Express

Although not as full featured as the full Adobe Photoshop application, Adobe Photoshop Express lets you edit your photos. This app supports

all of the basic photo-editing tools that you want, including cropping, rotating, brightening, and removing red eye. Live tiles display your recent photos. You can download the [Adobe Photoshop Express app](#) for free from the Windows Store.

③ Wikipedia

You probably see Wikipedia results in most Google (or Bing) searches, but the new Windows 8.1 Wikipedia app lets you search Wikipedia directly from your desktop. The app highlights featured images and articles. You can search for topics using the Search charm. One of the interesting things about the Wikipedia app is that it is 100 percent open source. You can get the [Wikipedia app](#) for free from the Windows Store.

② Evernote Touch

A handy app, but not quite up to the level of the OneNote app that's at the top of my list, Evernote Touch is the Windows 8.1 version of the popular Evernote note-taking app. Evernote Touch lets you record notes and photos and then organize them into multiple notebooks. Evernote Touch will sync notes across Windows 8.1/8, Windows Phone, Android, iPhone, and other devices. You can get the [Evernote Touch app](#) for free from the Windows Store.

① OneNote

The coolest new app for Windows 8.1 might be the OneNote app. Like its full Office counterpart, the OneNote app is a digital notebook that records notes, drawings, photos, and other digital content into multiple notebooks. The app supports the Share charm, allowing you to copy web content and information from other Windows 8.1 apps. There's a new camera-scanning optical character recognition (OCR) ability that lets you search for text in photos and screenshots. An insertion wheel lets you add a table, tag, photo, or list and also lets you paste to your notes. OneNote can sync across all your devices via SkyDrive. You can download the [OneNote app](#) for free from the Windows Store. ■

Introducing Windows Azure Backup

A quick and easy way to perform offsite backups

With so many Microsoft Azure solutions available, it's possible that one solution might have gone by unnoticed: Windows Azure Backup. As its name suggests, Windows Azure Backup allows Windows to make backups to Azure. No surprises there. However, what might surprise you is what's going on under the hood. Before I discuss how Windows Azure Backup works, though, I'll give you a high-level look at when you can use it and what you need to do so.

A High-Level Look

Windows Azure Backup lets you back up files in one of the following two scenarios:

- You can back up files from a standalone server.
- You can back up members of [Microsoft System Center Data Protection Manager](#) (DPM) protection groups. For this scenario, only DPM 2012 SP1 and DPM 2012 R2 can be used in conjunction with Windows Azure Backup.

If you want to make backups to Microsoft Azure, here's what you're going to need:

- An Azure account or subscription
- A certificate
- A passphrase
- The Windows Azure Backup agent

Let's take a closer look at these necessary components.



Robert Mitchell

is a senior support escalation engineer in the Windows Commercial Technical Support team at Microsoft, where he helps customers with Windows storage issues. He regularly posts to the [Ask the Core Team](#) blog.



Email



Blog

Azure Account or Subscription

You can create an Azure account on the [Windows Azure](#) website. Although I don't know how long the free offer will last, at the time of this writing, you can sign up for a free trial account. With a trial account, you can do some basic testing to see what solutions are right for you.

If you're already a service administrator as part of a larger account, you'll need to have Recovery Services added to your subscription. However, if you're creating your own account, you can add it on your own. For more information about the difference between an Azure account and an Azure subscription, see the [Manage Accounts, Subscriptions, and Administrative Roles](#) web page.

After you have your subscription in order, you need to create a backup vault. The vault is simply a storage construct used to hold your backups. You can put all your backups in a single vault or create multiple vaults to organize your backups.

Certificate

A certificate is used to grant access to the vault. The certificate must be uploaded to the vault. In addition, any server that you're going to back up to the vault will need to import this certificate.

You can use any valid SSL certificate issued by a Certification Authority (CA) trusted by Microsoft, whose root certificates are distributed through the Microsoft Root Certificate Program. The certificate must meet the following requirements:

- It must be an x.509 v3 certificate.
- There must be a .cer format file that contains the public key to upload to the vault.
- The key length should be at least 2048 bits.
- The certificate must have a valid ClientAuthentication extended key usage (EKU).
- The certificate should have a validity period that doesn't exceed three years.

- The certificate should reside in the Personal certificate store of the local computer on which you plan to install the Windows Azure Backup agent.
- The private key should be included during the installation of the certificate.

If you don't have a Certification Authority-issued certificate, you can use the MakeCert.exe tool to manually create a self-signed certificate. The MakeCert.exe tool is part of the [Windows Software Development Kit](#) (SDK). After you download the Windows SDK, you'll find the MakeCert.exe tool in the \Bin folder of the SDK's installation path. The following command will create a certificate that meets all of Microsoft's requirements:

```
Makecert.exe -r -pe -n "CN=AZUREBACKUP" -ss my  
-sr localmachine -eku 1.3.6.1.5.5.7.3.2 -e 12/12/2015  
-len 2048 AZUREBACKUP.cer
```

You might want to alter the expiration date, which is currently listed as 12/12/2015. However, don't set it to a date more than three years in the future.

Passphrase

The passphrase is used to encrypt the backups before they're copied into the vault. Selecting and storing your passphrase is important, because it's *not* shared with Microsoft. In other words, if you lose your passphrase, you won't be able to restore from the backups in the vault.

Although it's possible to use the same passphrase for all your servers, it would be akin to using the same password for every website you visit. Is it possible? Yes. Is it secure? No. It's recommended that you use a different passphrase for each server that you're backing up to Azure.

After you upload the certificate, keep the web interface to the vault open. You'll need this interface again shortly. Although I mentioned the passphrase, I haven't discussed how to set a passphrase yet. I'll do that when I'm discussing the Windows Azure Backup agent installation.

Windows Azure Backup Agent

The Windows Azure Backup agent can be downloaded from the [Windows Azure](#) website. The link is visible when you select the backup vault. You'll find that there are two flavors of this client:

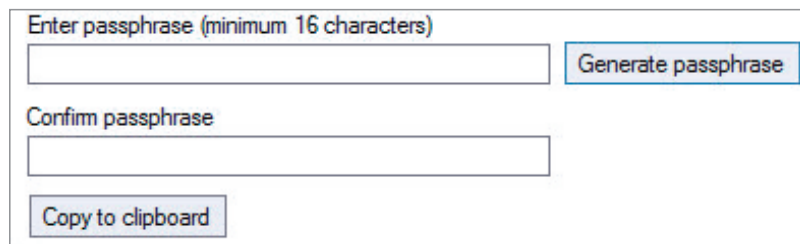
- [Windows Server 2012 R2](#), [Windows Server 2012](#), [Windows Server 2008 R2 SP1](#), and DPM
- Windows Server Essentials

Stepping through the installation is pretty standard. You'll be prompted for a cache location. This needs to be a volume with free space equal to or greater than 10 percent of the data set you plan to back up. This is where the Windows Azure Backup agent will put the backup while it's being compressed and encrypted, which occurs before it's copied to the vault.

You'll also be prompted for a passphrase. You can either create one or have the wizard generate it for you, as Figure 1 shows. Make sure you store your passphrase somewhere secure after you have entered it into the setup wizard.

When the agent installation is complete, you still need to register the server. This allows you to select the certificate you'll use on that

Figure 1
Creating the
Passphrase



The screenshot shows a web-based form for creating a passphrase. At the top, it says "Enter passphrase (minimum 16 characters)". Below this is a text input field. To the right of the input field is a button labeled "Generate passphrase". Below the first input field is the label "Confirm passphrase" followed by another text input field. At the bottom of the form is a button labeled "Copy to clipboard".

server and match it to the one you already uploaded to your backup vault.

For the standalone server scenario, you can register the server by running the Windows Azure Backup agent and selecting the Register Server option, as Figure 2 shows. When DPM is in the picture, you can register the DPM server from within the DPM Administrator Console, as shown in Figure 3. Both actions trigger the same registration wizard.




Figure 2

Registering a Standalone Server in the Windows Azure Backup Agent




Figure 3

Registering a DPM Server in the DPM Administrator Console

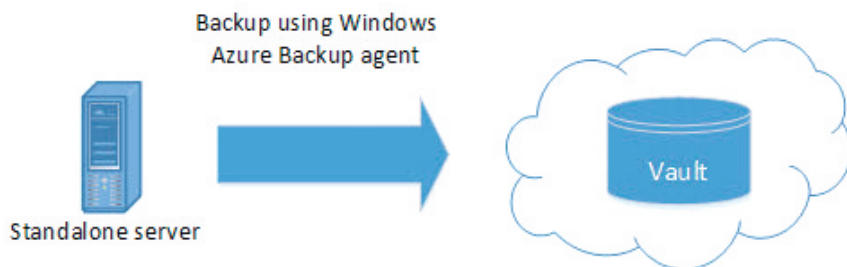
How Windows Azure Backup Works

So far, you've seen how to get Windows Azure Backup configured, but you need to know what it's actually doing when you use it. Because there are two different scenarios in which to use Windows Azure Backup, let's look at them individually.

Standalone server solution. As Figure 4 shows, the standalone server solution is straightforward. The Windows Azure Backup agent, which looks in appearance like Windows Server Backup, tells the Volume Shadow Copy Service (VSS) to create a snapshot of the data set to be backed up. For those of you unfamiliar with how VSS works under the hood, here's what typically happens: When a snapshot is taken, the data is put into a consistent state, then frozen (i.e., all changes are blocked) until the snapshot is created. Because you want

Figure 4

Understanding How
the Standalone Server
Solution Works



the data to be in a consistent state, writers will instruct VSS on how to handle their respective data sets. (Microsoft SQL Server has a writer, Microsoft Exchange has a writer, the registry has a writer, and so on.)

With Windows Azure Backup, VSS doesn't use any writers. For example, Figure 5 shows the VSSAdmin command being run to list snapshot information while the Windows Azure Backup agent is running. As you can see, no writers are being used.

```
C:\Windows\system32>vssadmin list shadows
vssadmin 1.1 - Volume Shadow Copy Service administrative command-line tool
(C) Copyright 2001-2013 Microsoft Corp.


Contents of shadow copy set ID: {b90be1f2-f669-4a4d-aed3-22e747d8a349}
  Contained 1 shadow copies at creation time: 1/22/2014 12:45:46 PM
    Shadow Copy ID: {5d44869a-e006-402d-8095-d58217348214}
      Original Volume: {C:}\\?\Volume{cf432842-6852-11e3-80b4-806e6f6e6963}\
      Shadow Copy Volume: \\?\GLOBALROOT\Device\HarddiskVolumeShadowCopy8
      Originating Machine: robtml.northamerica.corp.microsoft.com
      Service Machine: robtml.northamerica.corp.microsoft.com
      Provider: 'Microsoft Software Shadow Copy provider 1.0'
      Type: FileShareRollback
      Attribute: No writers, Differential
```

Figure 5

Running the
VSSAdmin Command
to Show That Writers
Aren't Being Used

Without a writer, data sets that need to be prepped for the freeze can't be prepped. The downside to all of this is that any data that requires a special VSS writer can't be backed up using Windows Azure Backup. This means no Exchange backups, no SQL Server backups, no system state backups—just offline data files. However, if you were looking forward to backing up SQL Server to your new vault, don't despair. You can use the DPM solution to do so.

DPM solution. DPM can protect various servers by backing them up to protection groups. Typically, that's the end of it. However, if you



also install the Windows Azure Backup agent on the DPM server, you can use it to back up members of the protection groups to your vault, as Figure 6 shows.

After the Windows Azure Backup agent is installed, the online protection option will be available in DPM, as you can see in Figure 7. Once this option is enabled and configured, your DPM protection groups will be backed up to your vault, extending the protection of your data to the [cloud](#).

Figure 6
Understanding How
the DPM Solution
Works

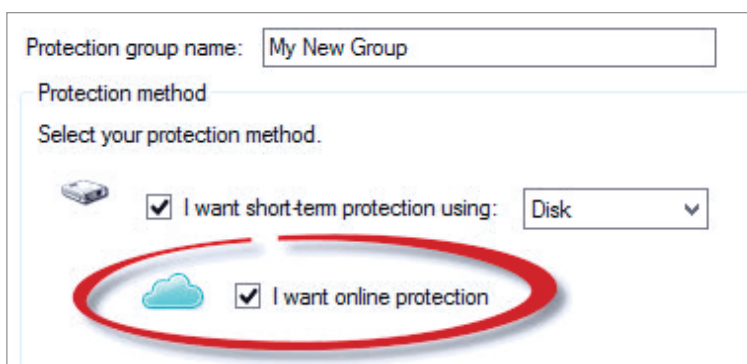


Figure 7
Choosing the Online
Protection Option in
DPM

In both the standalone server and DPM scenarios, the backup is created locally and encrypted using the passphrase that was supplied during the installation. Then the backup is copied to the appropriate vault, where it will remain, based on your retention policy.

Limitations

The standalone server and DPM solutions have different retention maximums. In the standalone server scenario, backups can be

Although it's possible to use the same passphrase for all your servers, it would be akin to using the same password for every website you visit. Is it possible? Yes. Is it secure? No.

retained in the vault for up to 30 days. This is configurable in the Windows Azure Backup agent's scheduling wizard.

In the DPM scenario, the maximum backup retention period is 120 days. In addition, there's a backup limit of either 120 backups or 850GB data size per protected data source. This can be configured in DPM's protection group wizard.

You should be aware of some other limitations as well. Windows Azure Backup can't be used when:

- A non-NTFS volume is used
- The drive type isn't fixed
- A volume is read-only
- A volume is offline
- A volume is on a network share

Quick and Easy Way to Perform Offsite Backups

Whether you're looking to back up files directly to the cloud or you're looking to strengthen the protection provided by DPM, Windows Azure Backup provides a quick and easy way to perform offsite backups. Just be aware that not all data sets can be protected by it. You can find additional tutorials and guides on the [Recovery Services](#) documentation web page. ■



Talking Exchange Server with Microsoft's Perry Clarke

Talking to any senior technologist can be a challenge for a journalist (which I'm not). The task of the journalist is to extract something juicy during the interview, while the interviewee has to resist all attempts to make himself open up and instead concentrate on passing the "message du jour." When I took over as HP's security lead in 2003, I received advice from the PR team that I should "always speak in paragraphs" because "you are less likely to be misinterpreted" and "whatever you say will be reported verbatim because it hangs together."

Moving forward 10 years, I find myself talking to Perry Clarke, corporate VP for Microsoft Exchange, and I'm the one attempting to get the big scoop. It's an interesting reversal, especially when faced with someone I've known for years. And Perry definitely doesn't answer questions in fully formed paragraphs. Instead, you get torrents of ideas and opinions, deep views, and reflections—all of which makes interesting listening, although it can be challenging to decipher.

Perry has spent his entire 17-year career at Microsoft working on Exchange Server and is now in charge of every piece of code written for Exchange, both on-premises and cloud (i.e., [Office 365](#)).



Tony Redmond

is a senior contributing editor for *Windows IT Pro* and the author of *Microsoft Exchange Server 2010 Inside Out* (Microsoft Press) and *Microsoft Exchange Server 2013 Inside Out: Mailbox and High Availability* (Microsoft Press).



Following Perry's recent "[Exchange Server: The Road Ahead](#)" post on the EHLO blog, I [canvassed the Exchange community for questions](#) and received a lot of feedback. I also checked out the [Microsoft Exchange Improvement Suggestions website](#). I gathered a lot of ideas for questions—certainly far too many to cover in an hour, especially in the cut-and-thrust of a debate between two people who care deeply about a particular technology and have devoted a fair portion of their careers to that technology. The topics Perry and I eventually managed to cover during our conversation include the following:

- The current state of Exchange
- The movement to the cloud and the future of on-premises software
- Mobility

I recorded the interview, and I later listened to the recording and wrote this article based on our conversation. Direct quotes from Perry are shown within quotation marks. The rest of the text is my interpretation of what Perry said (and my attempt to put it all into a logical order). Any errors in the article are entirely mine. I hope that you find the debate interesting.

The Current State of Exchange

We began our conversation with the current state of Exchange. I asked Perry if he's happy with where Microsoft is, in light of anecdotal evidence that customers aren't moving to [Exchange 2013](#) as quickly as they moved to previous versions, despite some major engineering advances—perhaps because of fears about poor quality, issues with interoperability, some missing functionality, and reports of poor server performance. I asked whether the current situation is similar to Exchange 2000, when customers held back because although a new architecture was introduced, there was also a lack of compelling new features—meaning that Exchange 5.5 was deemed to be “good enough.”

Although you'd expect any product's development supremo to defend his work, Perry surprised me by saying that he is happy with the Exchange product group's position, pointing out that they're in a much better position at this point in the development cycle (a year after shipping a major release) than they were at the equivalent time post-Exchange 2010. A common code base to span on-premises and cloud is in place and working, Exchange Online has taken on millions of new mailboxes, and the data reported from Microsoft's experience of running the service is being factored into decisions about the future development of Exchange. Perry pointed out that a flawed product couldn't possibly cope with the stresses and strains that Exchange Online undergoes. The Exchange 2013 code handles four times as many cloud mailboxes today as it did a year ago, and all of Microsoft's internal metrics report better performance and availability, as well as far fewer alerts than with Exchange 2010.

I demurred on this and pointed out that many of those who attempt to deploy Exchange 2013 on-premises have observed that it requires far more resources (mostly memory) than Exchange 2010 and seems to be more fragile. The loss of functionality such as S/MIME support is also being felt, as is the need to change the administrative model. I did concede that a huge amount of knowledge and experience now exist regarding Exchange 2007 and 2010 deployment, and that it will take time to accumulate the same expertise for Exchange 2013.

Perry also said that on-premises customers get a huge advantage from the fact that the "exact same" code base is now in use for both platforms and that [new code is stressed by being run in Office 365 for several weeks before it's released to on-premises customers in a cumulative update](#). "There isn't a part of the core Exchange Server that is not shipped to our on-premises customers. We don't have a special Exchange that runs in the cloud." Perry said that it should "comfort our on-premises customers that the code that they run is validated by running against millions of mailboxes before they have to deploy it."

I pointed out that Office 365 is a very structured environment that thoroughly exercises many parts of the product by running code at extreme scale but misses a lot of the detail found in customer deployments. For example, Office 365 doesn't permit server-side integrations with third-party products and insists on the use of the latest clients, whereas almost every on-premises deployment will use one or more third-party products alongside Exchange and involve a mix of clients. Perry agreed, but he noted that "the problem [of testing for specific customer environments] has always been there" and said that Microsoft "hasn't taken steps back on what we have done historically to address that issue."

We discussed interoperability with Exchange 2007, an area of functionality that isn't exercised or tested by Exchange Online and one that seems to be causing problems in customer projects. Perry acknowledged that "the n-1 release [Exchange 2007] has always been a challenge to get validation of the interoperability story." He noted that Microsoft still has some Exchange 2007 servers running production workloads and test labs but that customer environments have become "more complex," which makes interoperability testing more difficult. Given the nature of some of the infrastructures deployed into customer sites (including third-party software), it's impossible for Microsoft to test every aspect of interoperability, but the company is "not taking steps back" from the responsibility. Perry also said that the extensibility models (from Exchange 2010 onward) that now exist are much better and should prove to be more robust.

Exchange 2013 introduced a new problem in the need to continually increase the scale (of Exchange Online) and keep the code base moving. Perry said that "getting our cumulative update pipeline to be excellent has been an interesting challenge." Similar problems with quality updates had occurred with previous versions and had to be worked through then. A rigorous exercise was conducted during the past summer to figure out where the problems occurred in updates. Although better testing at scale was being done through deploying new

code into Exchange Online, regressions were still creeping through. An example of a recent regression in [Exchange 2013 CU3](#) is when [Windows XP clients can't use Internet Explorer 8 to connect to OWA](#).

In Perry's mind, a simple contract exists between the Exchange team and on-premises customers ("just apply the updates and your lives will get better"), so even a single regression is too much. The kind of regression seen in Exchange 2013 CU3 "shouldn't happen," and even though Microsoft feels confident that it has a handle on the problem and has deployed new checks and testing to improve the code that goes to on-premises customers, the company still needs to work hard to ensure that high-quality updates are delivered in the future. Perry acknowledged that customers won't believe the story of quality in Exchange 2013 until they see a stream of successful updates that demonstrate that fact. He said that internal Microsoft data shows that the company is in better shape to deliver updates than at any time in the past, but he also pointed out that it's difficult to implement the kind of functionality that's demanded by customers in older clients, especially when those clients don't support some of the basic underpinnings required to enable new functionality in components such as [OWA](#).

At no point did Perry attempt to say that everything is perfect with the new code. Indeed, he said that every new product creates some new "bumps in the road" for on-premises customers. But because most customers don't deploy until SP1 comes out, it gives Microsoft the chance to discover what bits are missing or don't work and to fix them in SP1. Perry pointed out that with Exchange 2010 there was a "bigger freak-out with what was missing in OWA [in the RTM release] than this time around" (OWA was rewritten for Exchange 2010 SP1). He said that there had been a huge reaction within and outside Microsoft when Exchange 2010 appeared because of some interoperability problems and other issues but that SP1 made everyone forget the original problems. In his view, Exchange 2013 is better than Exchange 2010 was, and we'll see the same degree of improvement when SP1 appears.

I asked why the same story has played out for many releases of Exchange, in which the RTM release is flawed and SP1 saves the day. Was it a case of Microsoft organizational politics requiring Exchange to ship at an arbitrary date, as in the case when Exchange 2013 RTM appeared alongside the other products in Office Wave 15? Perry vigorously disagreed with this hypothesis. So why did we see so many [problems in Exchange 2013 RTM](#)?

It's a complex issue. Perry pointed out that the "feedback loop" from customers to Microsoft is "too slow" because it takes on-premises customers a long time to deploy and test new software. The period between new features being designed, coded, tested, and shipped by Microsoft and when the product is eventually used by the "median customer" is "half a decade." No software development group can wait five years to hear whether they have done a good job.

Perry concludes that the only way of getting feedback at scale is therefore to put a product into the marketplace, which is what happened with Exchange 2013 RTM. The customers who wanted the feature set were able to deploy and use the software successfully, whereas those who needed some things to be fixed had to wait for a release like CU1. Many others will wait for SP1, and some will wait even longer. Perry concluded that this is the "nature of the on-premises game." He also said that "there is no customer on the face of the planet that is capable of doing a [major] Exchange deployment in less than three years [after the previous major release]. Systems are on a half-decade cycle: That's a fact of life."

Some might quibble at this, and the exception will prove the rule—but I think this statement expresses some of the natural frustration that software developers have when their work isn't used as quickly as they feel that it should be. Of course, the fact that Exchange has such a large installed base is both a benefit and a drag because it does take time to modernize the bulk of any base. This isn't a new phenomenon; the same is true for operating systems, and it was also true in the 1980s when the team I worked on at Digital struggled to convince customers

to deploy new versions of [ALL-IN-1](#), which was the leading corporate email system at the time. But Perry really values the size of the Exchange installed base because it's a huge asset to Microsoft.

Microsoft's new servicing model for Exchange 2013 is therefore based on two principles. The first is that Microsoft will "never go dark for two or three years," which means that the product will be in a state of constant improvement through updates shipped to customers on an incremental basis. The second is that some method has to be adopted whereby customers can keep pace with the development group by deploying the incremental releases. These principles have resulted in the quarterly cadence that now exists for Exchange 2013 cumulative updates.

In explaining the logic behind cumulative updates, each of which is a full-blown version of Exchange that can be deployed from scratch, Perry reflected that a traditional release is composed of three parts: incremental change to UI, improvements in existing functionality, and increased quality (bug fixes). Any update will have different proportions of the three parts: Some might change the UI dramatically (think of the OWA update in Exchange 2010 SP1); another might focus on increased quality; and another might roll out new features. Updates can also include architectural shifts to take advantage of new or emerging technologies and changes in hardware capabilities (the shift Exchange has made to trade memory for disk I/O is a good example of how Microsoft has used the fact that memory is getting dramatically cheaper to improve performance), as well as the evolution of what the user experience should be in light of new technology and the current environment (mobile computing is an example here).

Perry said that Microsoft's challenge is "how to stream the incremental wins into the update path and put the disruptive stuff into buckets that customers can consume." A major release changes the entire ecosystem around Exchange. Microsoft has to explain to its field and partners (major companies such as HP, as well as partners throughout the world that assist customers to deploy Microsoft

technologies) what the major changes are and why they're important, including the "deep thought" that has driven the change. Customers take that information and knowledge about the new software and put it into context with their own business requirements and existing infrastructure to figure out how they can build a deployment strategy that justifies the necessary expenditure for the migration project, including the capital expenditure for new computers. Perry said, "No one is asking us for the disruptive stuff at a faster cadence than we're delivering. If anything, they would like that pace to be slower."

So we end up with a [release cycle composed of multiple incremental updates leading up to a major release every three years or so](#). The shift to the new cadence has caused some problems to emerge that customers see as quality issues slipping through into the updates. Perry's view is that the model is now settling down as Microsoft and customers become accustomed to it, and he is confident that the "incremental stream [of updates] will be much more robust going forward."

Exchange and the Cloud

Getting back to the topic of the cloud and why so much importance is being put on cloud platforms today, Perry and I discussed how Microsoft developed the plan to transform Exchange into a cloud-capable application. Perry related that after the release of Exchange 2003, Microsoft had to chart out a plan to bring Exchange into a state where the team could deal with some of the technical developments that they could predict at the time and fix some of the deep technical and architectural flaws that had been introduced in Exchange 2000. Most of the flaws were addressed in Exchange 2003, but some have taken three releases to move Exchange from a product where components were "tightly coupled and loosely structured" to the situation that exists today where the focus is on "loosely coupled, tightly structured." (The best example of where this change has occurred is in the way that Exchange 2013 insists that [servers communicate using a small number of well-defined protocols](#) such as SMTP, Exchange Web Services, and MRSProxy.)

In 2005, the cloud was getting some serious traction and looked as if it would be a huge influence on the future. However, the perceived wisdom was that there was “no way that you could take a product team, code base, and knowledge base [that worked on-premises] and take it forward into the cloud.”

In some ways, the challenge of embracing the code paralleled that which faced Microsoft when the company developed Windows NT to break into the enterprise market and move away from its desktop client roots. It seemed like “a blank sheet of paper and a new organization would be required to move forward” to make applications such as Exchange cloud-capable. Perry said that two problems were obvious with such an approach. Internally, Microsoft would have to abandon a lot of code, people, and knowledge. In addition, the company would have to abandon the on-premises installed base (which at the time was well over 100 million mailboxes).

I asked if it didn’t seem like the on-premises base felt that they were being abandoned today, given Microsoft’s apparent eagerness to move everything to the cloud. Perry acknowledged that this perception might exist, but he pointed out that in 2005 Microsoft had three options:

- Ignore the cloud and concentrate on on-premises software
- Build a new team and a new product based on cloud technology
- Figure out how to accommodate both platforms

The first option would have completely ignored a platform that is now quite important and ceded command of the email cloud market to Google. The second option would have screwed the installed base and forced them to continue running software that gradually became more and more obsolete. This is a valid option for a newcomer to the email market and is the path taken by Google when the company developed Gmail. This option would also have caused “panic” within Microsoft as parts of the company struggled to retain customers—an influence that would have compromised the ability of the product team to deliver the

new product. The third option is the one that Microsoft took, which with the benefit of hindsight has proven to be the most valuable—even if the resulting transformation in the architecture and code base has created some problems that customers see as flaws and quality issues. Perry believes that Microsoft is getting a handle on these issues as the code becomes more settled and Microsoft gets better at leveraging the single code base to serve both on-premises and cloud platforms. Some evidence of this is seen in [Exchange 2013 CU3](#), which has had fewer problems reported than RTM, CU1, or CU2.

Perry said that there are few customers today who aren't thinking about whether they should move to the cloud. He did say that "a very large percentage of customers are *not* moving to the cloud" but that "two years ago when customers were thinking about what to do [for the next big upgrade], the percentage of customers who were considering a cloud strategy was much smaller." He believes that it takes time for companies to crystallize their thoughts around new developments and to figure out when it's the right time for them to adopt a new platform. The fact that Office 365 has delivered a robust and reliable service since June 2011 has helped many customers make up their minds.

If a company doesn't know what to do with the cloud, they can [create a test Office 365 domain to kick the tires](#) and figure out what's good and bad about cloud-based systems. It's very easy to run a cloud pilot, compared with the organizational and deployment effort required to run an on-premises pilot of new software. This ease of testing helps IT become more agile and make better decisions.

Given that many customers find it difficult to deploy new software as quickly as Microsoft releases it, Perry and I debated how much of the installed base would end up using Office 365—which means that Microsoft takes care of deploying and managing the software. Would we get to a situation, as predicted by some commentators, in which the complete installed base would move to the cloud, or would a substantial portion remain on-premises?

Perry's view is that even though there are many "perceived thought leaders" (including some inside Microsoft) who predict that 100 percent of customers will be in the cloud, this line of thinking is wrong and many customers have valid operational and business reasons to remain on-premises. Customers are hearing from people who are respected in the market (including market analysts) that everything is going to the cloud, but this doesn't make business sense. In Perry's words, "It's nuts."

Although Perry has been told for his entire 17-year Microsoft career that email is dead, he pointed to the "level of investment that customers are making to do email every fraction of every second of the day" to enable access to email everywhere from multiple devices. This creates "interesting challenges in data security, privacy, and data leakage." Microsoft has had to think through these problems in order to deliver Office 365 at scale, and that experience is, in Perry's view, "deeply valuable for our customers."

Given the [current run-rate for Office 365](#) and the available evidence in the market, Perry and I both agreed that somewhere in the region of 40 percent of the [total installed base will likely continue using on-premises software](#) for at least the next several years (in my view, probably much longer than that), which gives Microsoft a pretty big reason to ensure that this software remains highly functional. Perry acknowledged that this is actuality and pointed to the large investment that Microsoft is making to ensure that both on-premises and cloud platforms are supported with great interoperability between the two.

Mobile Computing

One of the big changes that Perry has noted over the past few years is the change in focus of customers from a simple cost per mailbox discussion to how they can make their users more productive. In the past, customers wouldn't talk to Microsoft about new software unless that software could reduce their cost; now they want to know what features exist that will make their employees more productive in

every sense of the word. The shift in IT results in a difference in the way that Microsoft thinks about products like Exchange. Part of that shift has been provoked by the explosion in the capability of mobile devices and the networks that allow people to remain perpetually connected to services such as email.

I asked about Exchange ActiveSync (EAS) and its future, because not much seems to have changed in the past few years. Perry made the bold assertion, “I don’t think that Apple would have been so successful with the iPhone against RIM in the enterprise if we hadn’t created the EAS ecosystem. It’s theoretically possible that Apple would have built a great client (for Exchange) using other protocols, but it’s unlikely that they would have invested the necessary resources on their own. There’s just no way that they [Apple] would have been able to engineer a MAPI client [to connect to Exchange].”

It’s certainly fair to say that [Exchange ActiveSync \(EAS\) is the de facto standard for mobile device connectivity to Exchange](#). Early versions of the iPhone were limited to POP3 and IMAP4 connections, but once Apple had licensed EAS, iPhones certainly proliferated inside large corporations. There’s no doubt that the advent of the iPhone, followed by Android and Windows Phone devices all using EAS to connect to Exchange, marked the start of RIM’s decline as fewer companies could justify paying for expensive BlackBerry Enterprise Server (BES) licenses when Microsoft provided EAS free from Exchange 2003 SP1 onward.

Is this a case of superb business acumen on the part of Apple to figure out that EAS was the best method of communication, or a case of blessed serendipity when everything came together at the right time to benefit Apple? Perry acknowledged that his statement was an extreme way of making the point but said that there is no question that the evolution of the EAS ecosystem and the standard that was created has helped Apple.

Moving forward, Perry said that the arrival of much more capable mobile devices has made it harder for platforms that use a

“sync-based approach” (i.e., ActiveSync) to keep up to date with the evolution of functionality on the server. People want to use the latest features on their new mobile devices, but device manufacturers can’t update their clients quickly enough—and the EAS protocol doesn’t support a lot of the features that are available to other Exchange clients. Being able to go to a repository such as the iOS app store and download code that exposes the full functionality of Exchange is a “dramatic game changer” that will drive user productivity. Moving the logic from clients to the server gives Microsoft the ability to deliver the “complete user experience” across the whole spectrum of mobile device form factors without the need for multiple vendors to create similar functionality numerous times. Another advantage is that Microsoft can take advantage of features that exist on each platform, such as the Safari browser on iOS. All of this comes together in a very real sense in [Outlook Web App for Devices](#), released for iOS in July 2013.

The natural question then arises regarding whether Microsoft is going to decommit from EAS and concentrate on apps that are published by Microsoft and refreshed on a frequent basis. Perry’s emphatic response was that although providing its own mobile clients will be “an increasing part of the story,” Microsoft continues to invest heavily in EAS, which has more clients connecting via the protocol than ever before, and he claimed that “in the order of a billion mobile devices ship with EAS on them annually.”

Perry also pointed out that Microsoft expends a lot of time to encourage EAS licensees to do the right thing when they [implement EAS inside their code](#) and that the Exchange group has also addressed some “mistakes” that have been made in EAS and improved the ability of the server to resist the effect of badly behaved clients. He noted that the sheer size of the EAS ecosystem means that the impact of any bug is “just huge” because it’s felt by literally hundreds of millions of users and that Microsoft takes the responsibility of delivering a highly reliable and robust EAS protocol very seriously.

At the same time, Microsoft will dedicate the necessary resources to satisfy the need for highly functional mobile clients that can take advantage of the latest browser technology running on mobile devices. Perry said, “The line has blurred when it comes to primary computing devices,” meaning that Microsoft must deliver a good experience on high-end smartphones and tablets. The discussion about how best to provide a great mobile experience continues on an ongoing basis with Apple, Samsung, and other vendors, as well as with Microsoft’s own Windows Phone development team. Perry noted that “it’s not just email clients anymore; there’s also the ability to use Exchange Web Services to interact with the server in interesting ways. The amount of innovation and investment that’s going on there is just huge. It’s a tremendously exciting space.”

My interpretation of what’s happening is that EAS will continue as the lowest common denominator for mobile device connectivity to Exchange (albeit capable of enabling highly functional clients) in tandem with a set of browser-based clients that expose more of the server functionality. Microsoft can push new releases of these clients out through app marketplaces far more quickly than they can get features enabled by mobile device vendors through EAS extensions. Given the variety (not always good) of EAS-enabled client applications implemented across Windows Phone, Android, and iOS platforms, I think this is a reasonable approach—but we’ll have to wait and see how OWA for Devices delivers over the next few years.

In Closing

Perry and I discussed the need for better email security and end-to-end encryption in light of the PRISM revelations and the requirement to protect the content of messages as they pass between companies. Perry said that the “silliness of SMTP being sent in the clear will be fixed” and that Microsoft will work with partners to resolve the problem in a standards-based way. Although no firm date can be put on when this will happen, it should be reasonably soon. No doubt we

will all have to upgrade our software (hopefully through an incremental release) to get this level of protection.

We finished up by talking about the upcoming [Microsoft Exchange Conference \(MEC\)](#), taking place in Austin, Texas, in early April. Perry is looking forward to MEC because he believes that the Exchange team will be able to lay down some big markers to disprove the oft-asserted wisdom of many so-called experts that email is going away sometime soon. He wants to be able to show that a server like Exchange can act as the fulcrum for many forms of natural and effective collaboration. It will be interesting to see how the MEC agenda evolves to fulfill the promise contained in his words. We shall see in time! ■



Windows IT Pro Store

eLearning Classes

eBooks

On-Demand Training

In-Person Training

Posters

Videos

Plus you can **RENEW** your subscription or
UPGRADE to VIP membership while
you're there!

Stop by the store today!

WindowsITPro

Windows Server 2012 File Classification Infrastructure

Automate the control and organization of sensitive data with the FCI feature



John Savill

is a Windows technical specialist, an 11-time MVP, and an MCSE for Private Cloud and Server Infrastructure 2012. He's a senior contributing editor to *Windows IT Pro* and his latest book is *Mastering Hyper-V 2012 R2 with System Center and Azure* (Wiley).

Email



Twitter



Website



Blog



Organizations store a huge amount of unstructured data: Microsoft Word documents, spreadsheets, images, data files from applications—the list goes on. Although applications such as SharePoint help organize such data, the reality is that only about 10 percent of data resides in apps like SharePoint. The rest sits on file servers with little control or management. Most organizations have no idea what data is on their file servers; it's just a mess.

Add to this the different requirements organizations have for data handling: restricting who can access data, ensuring it's encrypted, ensuring it can't be printed/copied/forwarded as part of Data Loss Prevention (DLP), ensuring it's backed up and kept for a certain duration—or conversely—making sure it's deleted after a certain duration.

The need for data classification and controls has never been more important than now. Regulatory compliance is getting more and more stringent. Organizations face massive fines and possible jail time for senior leaders who ignore compliance requirements. And then there's the huge loss of confidence in organizations that “lose” customer data.

Many companies try to classify data in a number of ways:

- Place documents in different locations
- Create backup rules for, or custom scripts that backup/delete, data of certain types/ages
- Apply Active Directory Rights Management Services (AD RMS) policies or encryption to sensitive data

- Use Windows BitLocker Drive Encryption to protect specific volumes containing important data

The problem is that many of these approaches rely on the information worker to make the correct decision in placing or classifying the data—a risk that organizations shouldn't take. A better option is to have the file server that houses the data scan it for social security numbers, credit card numbers, special project names, and the like and then automatically classify it. Once the data is classified, you can schedule tasks (e.g., back up, encrypt, rights protect, move, delete) based on the data's classification.

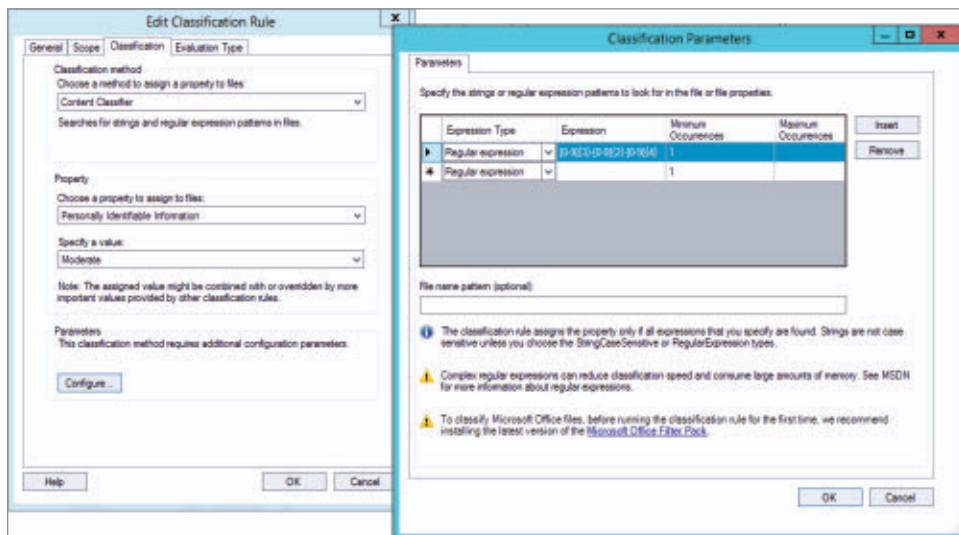
Meet File Classification Infrastructure

I remember creating a session for TechEd 2006 on a new Windows Server 2003 R2 feature called File Server Resource Manager (FSRM). The FSRM component brought capabilities to Windows file servers beyond the basic volume quota capability that was part of the base operating system.

FSRM lets you assign quotas to groups of users at the folder level. You also control not just how much space is used, but *how* the space is used by enabling real-time file screens that block specific file types. For example, you could create a file screen that blocks MP3 files. If a user tries to write an MP3 file, he or she gets an access denied message. These customizable actions are configurable (e.g., an action could be an email to the user explaining why they got the message and include a link to corporate policies). FSRM also has great reporting capabilities that identify how file server space is being used and by whom.

In Windows Server 2008 R2, FSRM got a new capability: File Classification Infrastructure (FCI). This feature uses rules to automatically assign specific properties to files and then performs tasks on those files based on the classification. For example, a classification rule might search for strings in the format of a Social Security Number—nnn-nn-nnnn ([0-9]{3}-[0-9]{2}-[0-9]{4})—and if one is found,

Figure 1
Using File Classification
Infrastructure to
Classify Files



assign the data a Personally Identifiable Information (PII) property of *Moderate*, as shown in Figure 1.

Once the data is classified, a file management task searches for data whose PII classification is set to Moderate or High and then applies an AD RMS policy that restricts how the data is used. Other actions, such as encryption or moving the data, also could be taken. Essentially, the FCI feature involves a two-step process:

1. Classify data using automated rules.
2. Perform tasks on data based on its classification.

A huge benefit of classification over the normal processes of searching data and then performing some immediate action is that actions *don't* have to be immediate. Data is classified periodically or as it's created, and then many sets of actions can be executed later based on the data's classification. This is a very powerful capability for organizations.

Although this technology was great, it wasn't widely adopted, even though organizations cried out for this type of feature. The reason for the lack of adoption was fairly simple: companies didn't know how to get started. Out of the box, FCI included no standard classification

properties, no standard classification rules, and no standard tasks to perform based on the non-existent, out-of-the-box classifications. This meant organizations first had to work out what classifications they needed—halting nearly every company in its tracks. Organizations spent months working out classifications, ended up with hundreds of possibilities, and then the project fizzled out and never happened.

To combat the lack of adoption, Microsoft released a Solution Accelerator (free download) for Windows Server 2008 R2 called the [Data Classification Toolkit \(available online\)](#). The toolkit includes a large number of classification properties, classification rules related to common compliance requirements, and tasks based on the classifications and focused on AD RMS policies. The toolkit provides customers a project base on which they can build. The Data Classification Toolkit only has 14 classification properties, but these facilitate the handling of nearly all classification and compliance requirements. I outlined the toolkit's base properties in Table 1.

Table 1: Classification Properties Included in the Data Classification Toolkit

Classification Area	Classification Property	Possible Values
Information Privacy	1. Personally Identifiable Information (PII)	High; Moderate; Low; Public; Not PII
	2. Protected Health Information	High; Moderate; Low
Information Security	3. Confidentiality	High; Moderate; Low
	4. Required Clearance	Restricted; Internal Use; Public
Legal	5. Compliance	SOX; PCI; HIPAA and many more
	6. Discoverability	Privileged; Hold
	7. Immutable	Yes/No
	8. Intellectual Property	Copyright; Trade Secret and more
Records Management	9. Retention	Long-term; Mid-term; Short-term; Indefinite
	10. Retention Start Date	<Date>
Organizational	11. Impact	High; Moderate; Low
	12. Department	<Department>
	13. Project	<Project>
	14. Personal Use	Yes/No

Note that you can use the File Server Resource Manager UI or Windows [PowerShell](#) to look at possible values. For example, I can use the Server 2012 PowerShell commands in Listing 1 to look at the Data Classification Toolkit's values for Compliance.

Listing 1: PowerShell Commands to View the Data Classification Toolkit's Values for Compliance

```
$propertyDefinition = get-fsrmclassificationpropertydefinition Compliance_MS
Foreach ($possiblevalue in $propertyDefinition.PossibleValue)
{
    $possibleValue
}
```

```
Description      :
DisplayName       : PCI DSS
Id               : 2DD2F3EE-3BAB-45fc-B33F-65119B3B3C66
Name             : PCI DSS
PSComputerName   :
```

```
Description      :
DisplayName       : HIPAA/HITECH
Id               : A9E2C599-7DC4-4bf1-90DA-E949EF25D045
Name             : HIPAA/HITECH
PSComputerName   :
```

```
Description      :
DisplayName       : SOX
Id               : 0424473A-B85A-4071-8A8F-AB3F230864A0
Name             : SOX
PSComputerName   :
```

```
....
```

Fast Forward to Windows Server 2012 FCI

So what changed in [Windows Server 2012 FCI](#)? A lot has changed—and not just the FCI feature itself, but also how classification is now used.

FCI is still part of the FSRM role, which itself is part of the File and Storage Services role. That means to enable FCI, you must first install the FSRM role through Server Manager (\File and Storage Services\File and iSCSI Services\File Server Resource Manager) or PowerShell:

```
Install-WindowsFeature FS-Resource-Manager
```

A major change in Server 2012 FCI is how you manage classification properties. In Server 2008 R2, the classification properties are local to each file server, which means you must be careful to ensure the same classification properties are available on all file servers, or classifications could get lost when files are moved between file servers. This is typically achieved by maintaining a master/staging file server where all classifications, rules, and tasks are defined and then exporting the configuration to other file servers (in fact, a master/staging file server is still good practice in Server 2012).

In Server 2012, the classification properties have moved into a new container in an Active Directory (AD) forest's Configuration partition (\Services\Claims Configuration\Resource Properties). You can still add local classification properties to a server, but to use AD classification, you must run the Server 2012 forest preparation step, update the forest schema, and create new containers. The plus side is that classification properties are centralized and standard across all file servers.

Manage classification properties. You manage classification properties in the Active Directory Administrative Center (ADAC). All classification properties are disabled by default, so you must enable those that you want to use. Additionally, some classification properties (such as Company, Department, and Project) require values before they can be used. To manage resource properties, launch ADAC and navigate to Dynamic Access Control, Resource Properties, where you can modify and enable properties and create additional classifications. Figure 2 shows the built-in Resource Properties in Server 2012. Note that in addition to the 14 classification properties from the Data Classification

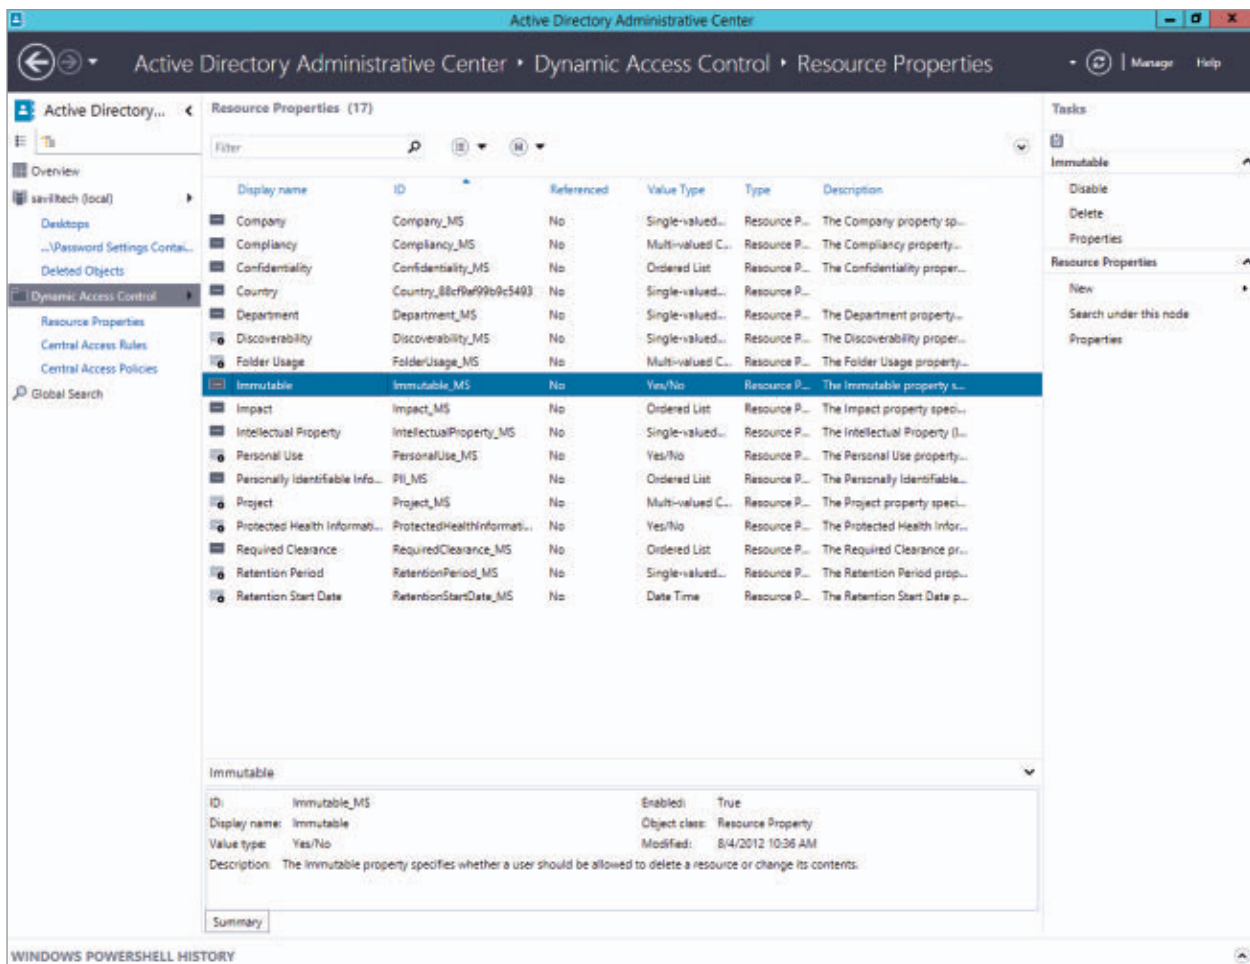


Figure 2
Managing
Classification
(Resource) Properties
with Active Directory
Administrative Center

Toolkit for Server 2008 R2, there are some additional ones; namely, Company, Country, and Folder Usage.

Management of classification properties is accomplished within ADAC's Dynamic Access Control, a major new feature in Server 2012 that lets you use data classification to control resource access. Dynamic Access Control is beyond the scope of this article, but at a high level it lets you control access to resources based on classification data and attributes of the user and machine trying to access the data. For example, you could use Dynamic Access Control to grant a

level of access if the department of the user matches the department classification of the data, avoiding the need to maintain hundreds—if not thousands—of groups just for access control. This means data classification isn't intended just to help secure and organize data for compliance, but also to manage resource access in a far more auditable and logical way than using ACLs on every file.

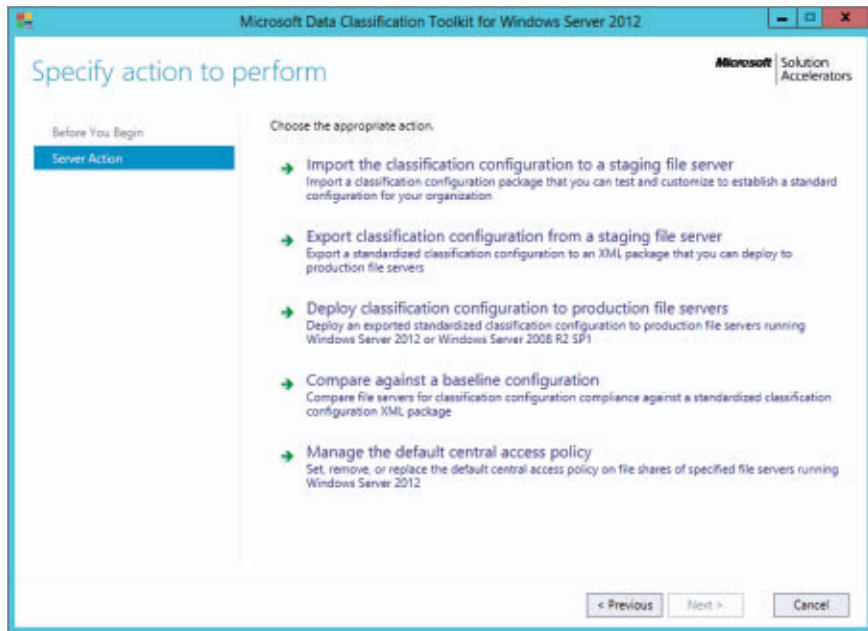
Get classification rules and management tasks. Shifting the focus back to FCI, the centralization and inclusion of default classification properties will certainly help organizations get started, but what about classification rules and file management tasks—are they standard in the box? The answer is no; however, Microsoft has updated the Data Classification Toolkit to work with Server 2012, which means you should download and run it. The toolkit will create a number of classification rules and file management tasks on your server.

Before you run the Data Classification Toolkit import wizard (which populates the targeted FSRM server), it's important that you enable most of the classification properties in AD and then refresh the FSRM Classification Properties. If you don't take these steps first, the import process will fail because the classification properties won't be available to the templates. Additionally, you should have AD RMS deployed in your organization so that you can apply a default AD RMS policy (that you can change later) and specify XML configuration files when you import the toolkit.

Configure classifications. The Data Classification Toolkit wizard (Figure 3) steps you through the process of configuring classifications. First, import the baseline, in-box classifications to a staging server. The toolkit encourages you to manage FCI on a single server, configure rules and tasks, and then export that configuration to all file servers in your environment. Once you've tweaked the configuration, export it to an XML file and deploy it to your production file servers.

The Data Classification Toolkit contains three baseline classification templates to help you get started. These are just XML files that

Figure 3
Primary Actions of
the Data Classification
Toolkit



you can apply to your file server; however, some customization will likely be required. The classification templates are:

- Data Classification Toolkit Package.xml—A standard set of rules and tasks primarily focused on finding SSN and credit card information in data
- NIST SP 800-53 Classification Package Example.xml—NIST SP 800-53
- PCI-DSS Classification Package Example.xml—PCI-DSS

Note that you don't have to use the Data Classification Toolkit to import and export a classification configuration, but its templates provide a great starting point for your own rules and tasks. Alternatively, you can use PowerShell.

To manage classification beyond the global classification properties stored in AD, use the FSRM tool via the Classification Management and File Management Tasks navigation nodes. Figure 4 shows the classification properties available. Note that the scope is visible

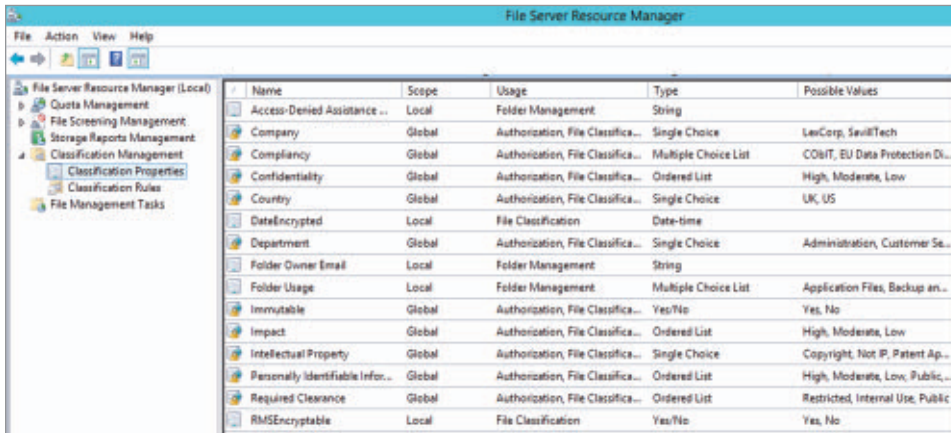


Figure 4
File Server Resource
Manager Displays
Classification
Properties Available

for each classification property. The properties with a global scope were retrieved from AD. Note also each property has a usage type that shows whether it can be used as part of Dynamic Access Control authorization, FCI, and folder management.

The Classification Rules navigation node enables you to manage the rules that populate the classification properties. If you leveraged the Data Classification Toolkit and imported the standard Data Classification Toolkit Package.xml, you have a number of classification rules available; however, they are disabled by default (imported file management tasks also are disabled by default).

Take some time to look at the classification rules and how they work and create your own if necessary. Make sure you enable the rules you want to use in your environment, as shown in Figure 5. Look at the detailed properties of a classification rule. Note that you can also use PowerShell to ascertain classification properties, which gives you limitless flexibility. One action you might want to take for each rule is to customize the folders to which a rule applies (on the scope tab of a rule's properties) and how the values are set for classification properties.

Schedule classification activities. Your configuration now has classification properties and rules to set them. Next you need to tell FCI when it should perform the classification. Click the Configure

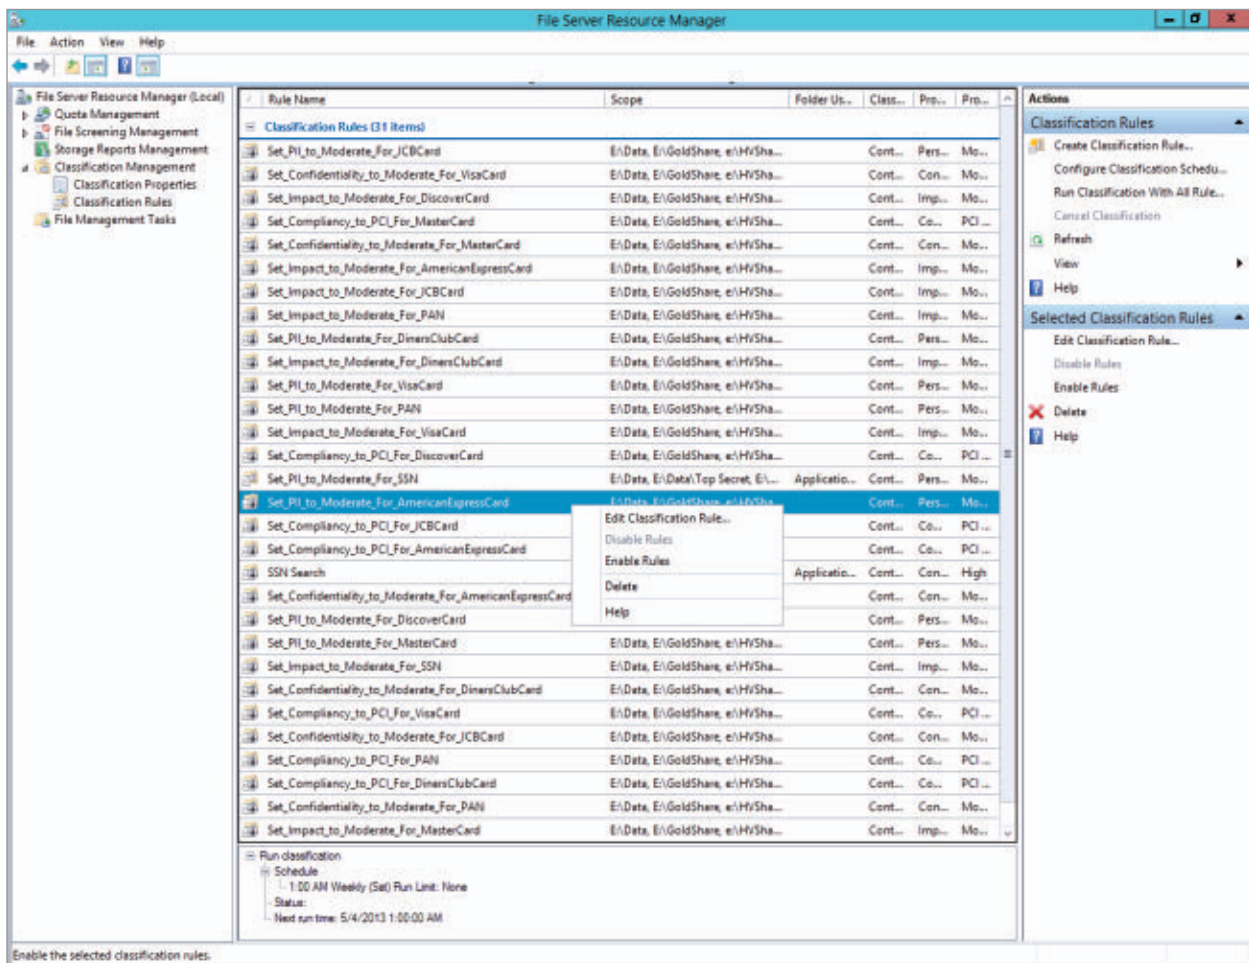


Figure 5

Enable a Classification Rule with the Enable Rules Action

Classification Schedule action, which opens the FSRM properties on the Automatic Classification tab. This lets you schedule scans to classify unclassified data. Another option, *Allow continuous classification for new files*, is a great new feature in Server 2012 that classifies data as soon as it's created.

Store classification data. At this point your data is classified, and a common question is: Where is the classification data stored? Classification data is stored in an NTFS alternate data stream in the file or folder that has the classification:


```
PS E:\unsc> Get-Item .\master_chief_eyes.jpg -Stream *
    FileName: E:\unsc\master_chief_eyes.jpg
Stream                      Length
-----
: $DATA                     39060
FSRM{ef88c031-595...       144
```

This means the classification is kept, even when you move the data between NTFS volumes. If you use Server 2012's new Resilient File System (ReFS), classification won't work because ReFS doesn't support alternate data streams in Server 2012. Additionally, if the data type supports it, the classification is stored within the document via the classification storage module. Microsoft Office is the primary application suite that allows classification to be stored in the application data, which also means the classification travels with the documents if they are stored in SharePoint. Other vendors could add support for classification storage in document data if they were so inclined. In Server 2012, the classification also is stored in the security descriptor to enable Dynamic Access Control authorization based on classification.

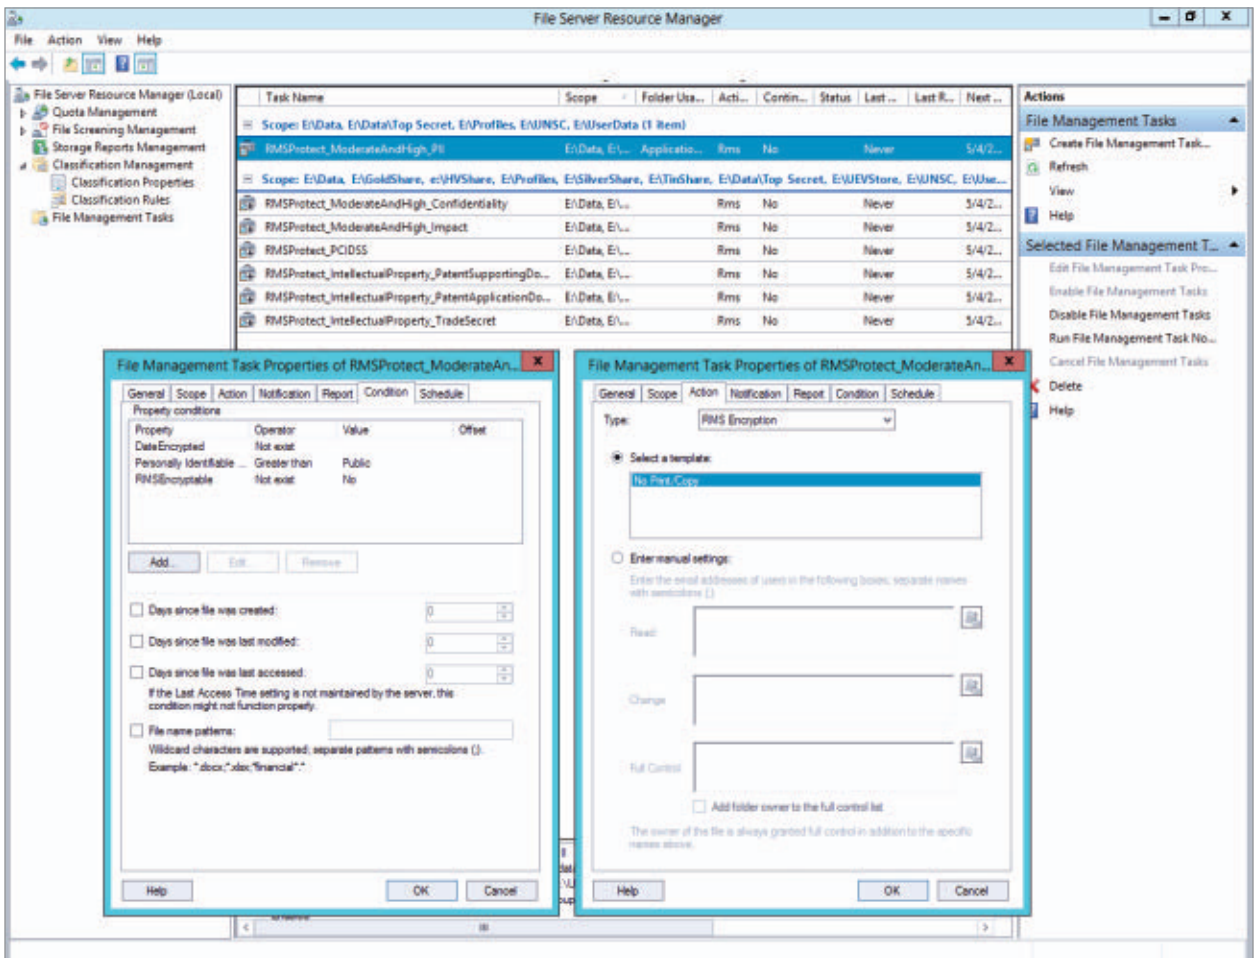
Use Windows Explorer. Windows Server 2012 and Windows 8 Windows Explorer exposes classification as a new Classification tab when you look at the properties of a file or folder, allowing direct manipulation of the classification; however, manually setting classification using Windows Explorer isn't practical for all organization data. One nice capability, though, is you can set classification at a folder level and then all folders and files in that folder will inherit that setting.

Perform management tasks. Although classifying the data is a huge step, you also want FCI to perform tasks based on those classifications. Use FSRM's File Management Tasks to create tasks that perform actions based on classification. If you used the Data Classification Toolkit to import a baseline configuration, some tasks are already configured, but they are also disabled. Look at the tasks and customize and enable the ones you want based on the classifications

for which you have rules. For example, if you enabled rules to set a value for the Personally Identifiable Information (PII) classification property, then you should enable tasks to perform actions based on that classification property. The included RMSProtect_ModerateAndHigh_PII (Figure 6), for example, performs a task if PII is greater than Public and then sets AD RMS policy on the data. Note that a custom type of action is available, which lets you perform almost any action provided there's a command-line method of doing so. Each task has its own schedule, and (like classification rules) the

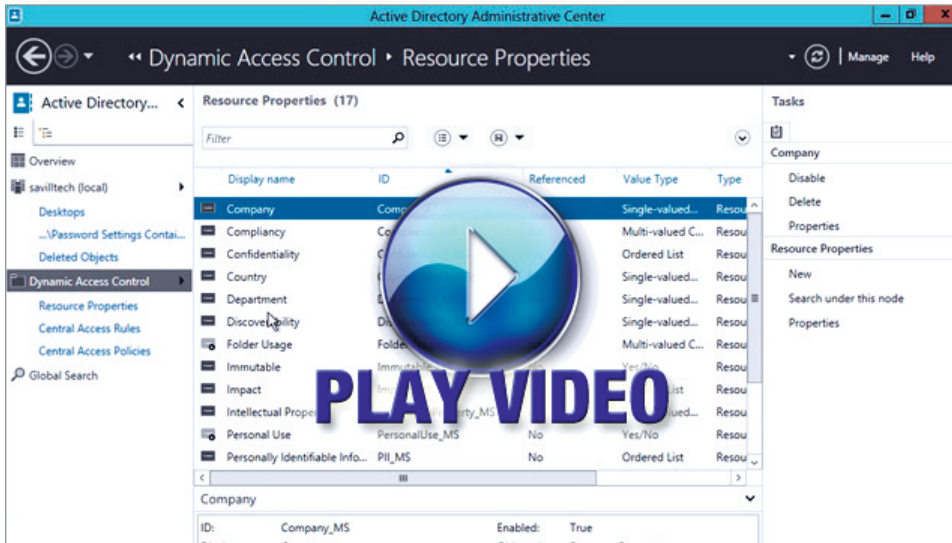
Figure 6

A File Management Task Protects Data Classified as PII Moderate or High



option to run continuously on new files can trigger tasks as classifications are applied.

All that's left now is to create some data files that contain the elements your rules search for, and your data is automatically classified. In the accompanying video, I walk through the major steps of using FCI.



Video

John Savill demonstrates the Windows Server 2012 FCI feature

Save with File Classification Infrastructure

FCI is a great technology, but before it can become useful to your organization—even before you implement it—you should understand what compliance levels your organization requires. Is your company required to adhere to certain regulatory standards? Does your company have its own standards for data retention, protection, and organization? Take time to understand these requirements, and then start implementing FCI, which comprises the following high-level steps based on a Server 2012 implementation:

1. Enable the AD classification properties you need to use and define values if necessary.
2. Install the FSRM role on file servers and the Data Classification Toolkit on one staging/master file server.

3. Import a template as a starting point. Customize the classification rules and tasks to your organization's exact needs.
4. Export the classification configuration from the staging/master file server to all other file servers.
5. Look at using existing classifications for other purposes such as Dynamic Access Control.

The in-box FCI implementation could save your organization a lot of money over third-party solutions. It also means data needs to be “local” to Windows file servers (i.e., direct-attached or mounted volumes from a SAN). FCI doesn't work for remote data such as that accessed over Server Message Block protocol.

I really only touched on the high-level capabilities of FCI. It's one of those technologies that could literally change the way an organization works, but it's also not widely understood or used. I strongly encourage every organization to look at FCI and discover how it might help organize and control unstructured data. ■

Exploring PowerShell's Group-Object Cmdlet

How to group objects based on existing and custom properties

One reason why [Windows PowerShell](#) is so flexible and maybe even a little harder to learn is that it doesn't have monolithic commands that do six different things. Instead, there are simple cmdlets that you can string together in a pipelined expression. Each cmdlet is designed for a single purpose. One cmdlet that you'll use often is Group-Object, which has the commonly used alias of *Group*.

As the name suggests, the Group-Object cmdlet puts objects into groups based on a property. This can be an existing property or a custom property. I'll show you how to group objects using both types of properties. I'll also show you some special techniques such as viewing only the total count of grouped objects and creating a grouped hash table.

Using Existing Object Properties

Typically, an existing property is used to group objects. When you use Group-Object, it writes a new object to the pipeline. Take, for example, the command:

```
Get-Service | Group Status
```

Even though this command starts with services, at the end of the pipeline is a `Microsoft.PowerShell.Commands.GroupInfo` object. You can see most of its properties in the sample output in Figure 1. The



Jeffery Hicks

is a Windows PowerShell MVP with almost 20 years of IT experience. He works as an independent consultant, trainer, and author. His latest book, with Don Jones and Richard Siddaway, is *PowerShell in Depth: An administrator's guide* (Manning, 2013).



Group property is a collection of the underlying objects that share the same property value—that is, all running or stopped services. The Name property reflects the name of each group. The Count property reflects the number of objects in each group.

Figure 1

Grouping Services by
the Status Property

```
PS C:\> Get-Service | Group Status
```

Count	Name	Group
94	Stopped	{AeLookupSvc, ALG, AllUserInstallAgent, AppI...
65	Running	{Appinfo, AudioEndpointBuilder, Audiosrv, BF...

Because the output from Group-Object is another object, you can use it like anything else in PowerShell. For example, consider the command:

```
Get-Command | Group Verb | Sort Count -Descending |  
Select -First 5 Name,Count | Format-Table -Auto
```

This command gets information about the PowerShell cmdlets, groups this information by the Verb property, sorts the grouped information by the Count property in descending order, and selects the first five. Figure 2 shows sample results. Perhaps this example isn't the most compelling, but I think it adequately demonstrates using Group-Object in a pipelined expression.

Figure 2

Grouping Information
by the Verb Property

```
PS C:\> Get-Command | Group Verb | Sort Count -Descending |  
Select -First 5 Name,Count | Format-Table -Auto
```

Name	Count
Get	412
Set	172
New	125
Remove	96
Add	84

Most of the time, administrators group on a single property, but it's possible to group on multiple properties. The assumption is that

you have objects that share a set of properties. The following code provides one example:

```
Get-Eventlog System -Newest 250 | Sort Source |
  Group EntryType,Source | Out-GridView -OutputMode Single |
  Select -ExpandProperty Group |
  Format-Table -GroupBy Source -Property TimeGenerated,
  Message -Wrap
```

This command starts by retrieving the 250 latest entries from the System event log. The entries are sorted by the Source property. The sorted results are then grouped first by the EntryType property, then by the Source property. The results are piped to Out-GridView, as shown in Figure 3.

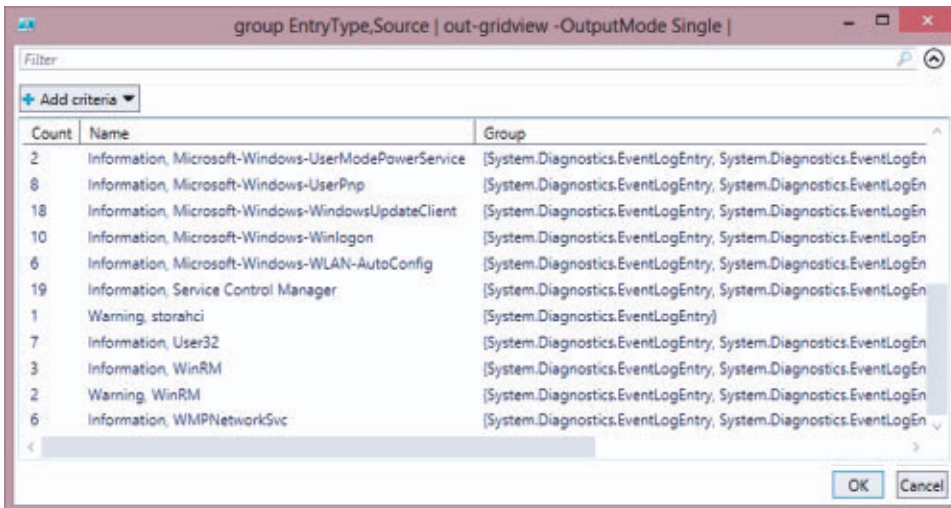
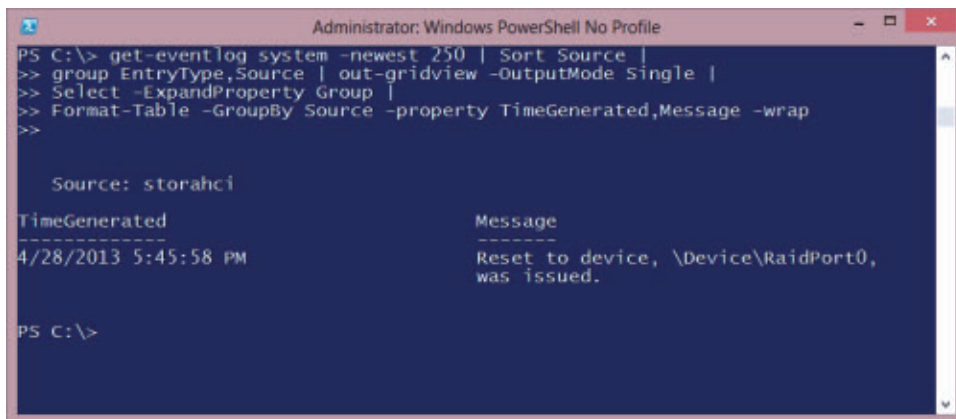


Figure 3

Grouping System Event Log Entries by the EntryType and Source Properties

The rest of the command executes after I select an entry and click OK. The command then expands the Group property, which is the collection of event log entries, and formats the results, as Figure 4 shows. This is possible because Out-GridView writes objects to the pipeline in PowerShell 3.0.

Figure 4
Expanding the Group
Property



```

Administrator: Windows PowerShell No Profile
PS C:\> get-eventlog system -newest 250 | Sort Source |
>> group EntryType,Source | out-gridview -OutputMode Single |
>> Select -ExpandProperty Group |
>> Format-Table -GroupBy Source -property TimeGenerated,Message -wrap
>>

Source: storahci
-----
TimeGenerated      Message
-----
4/28/2013 5:45:58 PM Reset to device, \Device\RaidPort0,
was issued.

PS C:\>

```

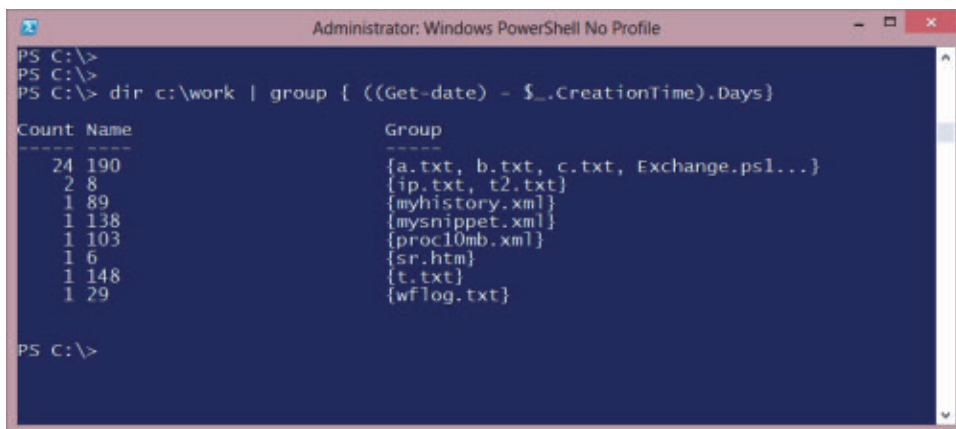
Using Custom Properties

You aren't limited to using existing object properties with Group-Object. You can group objects based on a value derived from a script block. Take, for example, the command:

```
Dir C:\Work | Group { ((Get-date) - $_.CreationTime).Days }
```

The `$_` represents each object in the pipeline. In this example, I'm retrieving all the files in the C:\Work directory and grouping them based on a calculated value that's the total number of days since the file was created. Figure 5 shows sample results.

Figure 5
Grouping Objects
Based on a Value
Derived from a Script
Block



```

Administrator: Windows PowerShell No Profile
PS C:\>
PS C:\>
PS C:\> dir c:\work | group { ((Get-date) - $_.CreationTime).Days }

Count Name                                     Group
-----
24 190 {a.txt, b.txt, c.txt, Exchange.ps1...}
2 8 {ip.txt, t2.txt}
1 89 {myhistory.xml}
1 138 {mysnippet.xml}
1 103 {proc10mb.xml}
1 6 {sr.htm}
1 148 {t.txt}
1 29 {wflog.txt}

PS C:\>

```

Viewing Only the Group Totals

Sometimes you might not care about the grouped results and only want to see the group totals. In this situation, you can tell PowerShell to skip the individual objects. For example, when all I want to see is a distribution of file types in my Scripts folder, I run the command:

```
Dir C:\Scripts -File -Recurse | Group Extension -NoElement |
    Sort Count -Descending
```

As you can see in Figure 6, I write a lot of PowerShell scripts. Using this technique is a terrific way to slice and dice data.

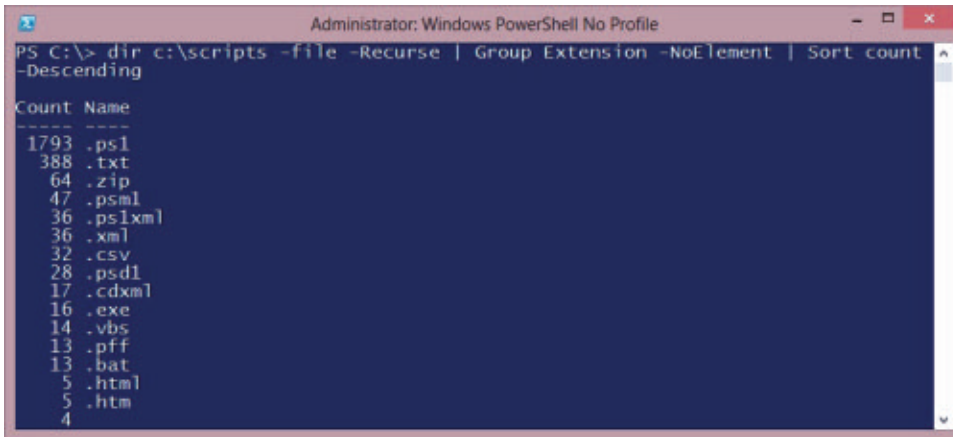


Figure 6

Viewing Only the Group Totals

Creating a Grouped Hash Table

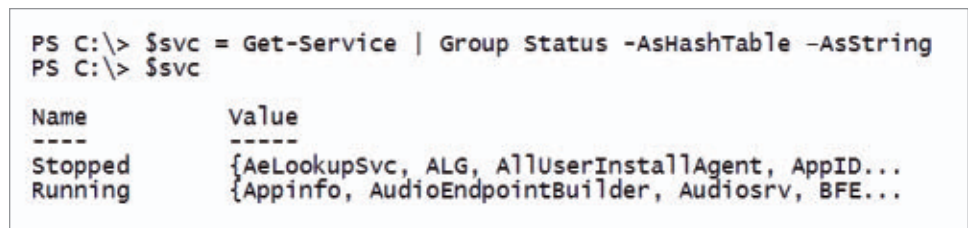
There might be situations in which you want to work with grouped data in a more interactive fashion. The easy way to accomplish this task is to turn the GroupInfo object into a hash table. A hash table, also known as an associative array (or a Dictionary object in the VBScript days), consists of a key/value pair. When you turn the GroupInfo object into a hash table, the Name property becomes the hash table key and the value is the collection of grouped objects. If you use this technique, I recommend that you filter out any objects that might result in a blank value.

To turn the GroupInfo object into a hash table and view its contents, you use commands like this:

```
$svc = Get-Service | Group Status -AsHashTable -AsString
$svc
```

Note the use of the -AsString parameter in the first command. Many object properties look like strings but are actually numeric values or enumerations under the hood. If you don't use the -AsString parameter, it will be very difficult to work with the hash table. It's difficult to know which properties you need to treat as a string, so I recommend always using it when creating a hash table. Figure 7 shows sample results from these commands.

Figure 7
Creating the \$svc Hash
Table and Viewing Its
Contents



```
PS C:\> $svc = Get-Service | Group Status -AsHashTable -AsString
PS C:\> $svc
```

Name	Value
Stopped	{AeLookupSvc, ALG, AllUserInstallAgent, AppID...
Running	{Appinfo, AudioEndpointBuilder, Audiosrv, BFE...

After you've created the hash table, you can use it for a wide variety of tasks. For example, if you want to see all the stopped services, you run the command:

```
$svc.Stopped
```

Figure 8 shows sample results.

Here's another example of creating a grouped hash table and viewing its contents:

```
$events = Get-Eventlog System -Newest 500 |
    Group Source -AsHashTable -AsString
$events
```

This command obtains the last 500 entries written to the System event log and creates a hash table based on the event log entry's Source property. You can see sample results from these commands in Figure 9.

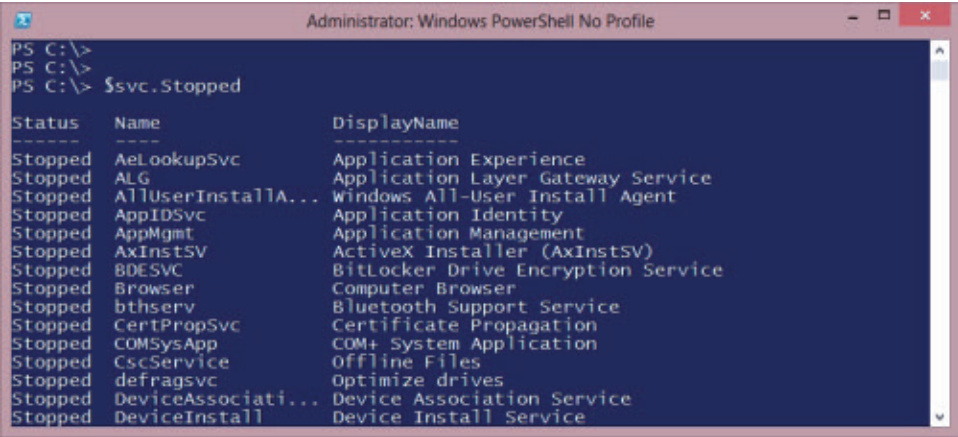


Figure 8
Using the \$svc Hash Table to See All the Stopped Services

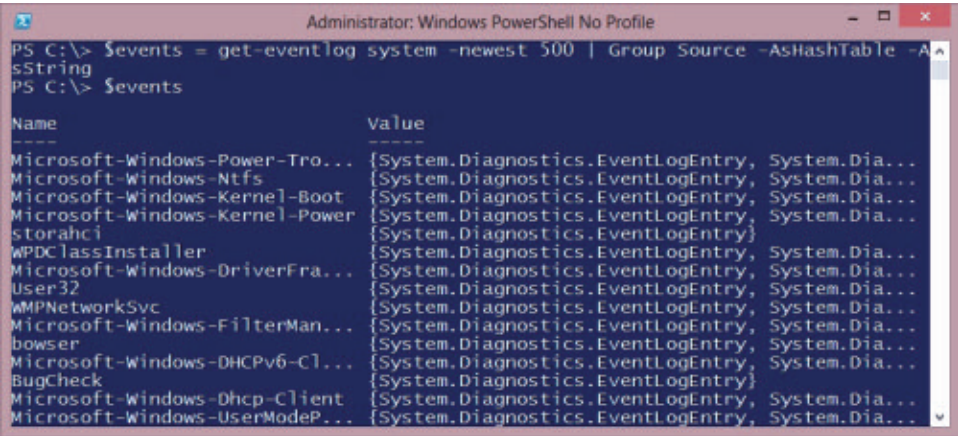


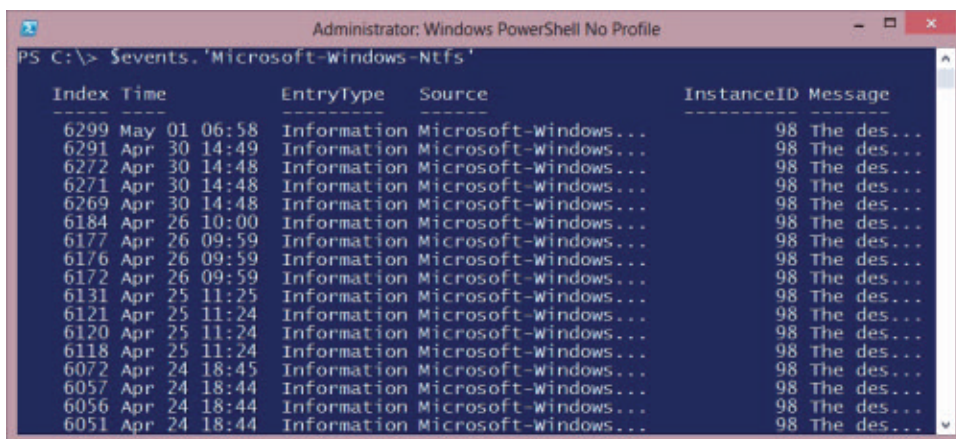
Figure 9
Creating the \$events Hash Table and Viewing Its Contents

You now have an object, \$events, that you can work with interactively to easily explore events from different sources. The handy thing about a hash table is that you can reference the value by treating the key as a property. (This is where tab completion is very useful.) Consider the following example:

```
$events.'Microsoft-Windows-Ntfs'
```

This command accesses an event source key (Microsoft-Windows-Ntfs) and displays all the corresponding event log entries, as Figure 10 shows. Even though the properties were turned into strings, some names have irregular characters and need to be quoted, as shown here.

Figure 10
Treating a Hash Table
Key as a Property



Administrator: Windows PowerShell No Profile

```
PS C:\> $events.'Microsoft-Windows-Ntfs'
```

Index	Time	EntryType	Source	InstanceID	Message
6299	May 01 06:58	Information	Microsoft-Windows...	98	The des...
6291	Apr 30 14:49	Information	Microsoft-Windows...	98	The des...
6272	Apr 30 14:48	Information	Microsoft-Windows...	98	The des...
6271	Apr 30 14:48	Information	Microsoft-Windows...	98	The des...
6269	Apr 30 14:48	Information	Microsoft-Windows...	98	The des...
6184	Apr 26 10:00	Information	Microsoft-Windows...	98	The des...
6177	Apr 26 09:59	Information	Microsoft-Windows...	98	The des...
6176	Apr 26 09:59	Information	Microsoft-Windows...	98	The des...
6172	Apr 26 09:59	Information	Microsoft-Windows...	98	The des...
6131	Apr 25 11:25	Information	Microsoft-Windows...	98	The des...
6121	Apr 25 11:24	Information	Microsoft-Windows...	98	The des...
6120	Apr 25 11:24	Information	Microsoft-Windows...	98	The des...
6118	Apr 25 11:24	Information	Microsoft-Windows...	98	The des...
6072	Apr 24 18:45	Information	Microsoft-Windows...	98	The des...
6057	Apr 24 18:44	Information	Microsoft-Windows...	98	The des...
6056	Apr 24 18:44	Information	Microsoft-Windows...	98	The des...
6051	Apr 24 18:44	Information	Microsoft-Windows...	98	The des...

Providing a Practical Example

To finish exploring Group-Object, let's look at a practical example. Listing 1 shows a script, FileExtensionAgeHTML, that analyzes a folder and creates an HTML report to group files by their extension and age.

Listing 1: The FileExtensionAgeHTML Script

```
$head = @'
<style>
body { background-color:#FFFFFFF;
      font-family:Tahoma;
      font-size:10pt; }
td, th { border:1px solid black;
         border-collapse:collapse; }
th { color:white;
     background-color:black; }
table, tr, td, th { padding: 2px; margin: 0px }
```

Listing 1: *continued*

```
tr:nth-child(odd) {background-color: lightgray}
table { margin-left:50px; }
</style>
'@
```

```
A $path = "C:\scripts"
   $files = DIR $path -Recurse -File
   $groupExt = $files | where {$_.extension} |
   Group-Object {$_.Extension.Substring(1)}

B # Create aging fragments.
   $30days = $files |
   where { $_.LastWritetime -ge (Get-Date).AddDays(-30) } |
   Group-Object {if ($_.extension) {
       $_.Extension.Substring(1)}} | Select Name,Count,
   @{Name="Size";Expression={
       ($_.Group | measure-object Length -sum).sum}} |
   Sort Count -Descending |
   ConvertTo-HTML -Fragment -PreContent "<h2>30 Days</h2>"

   $90days = $files |
   where { $_.LastWritetime -le (Get-Date).AddDays(-30) -and
       $_.LastWritetime -ge (Get-Date).AddDays(-90) } |
   Group-object {if ($_.extension) {
       $_.Extension.Substring(1)}} | Select Name,Count,
   @{Name="Size";Expression={
       ($_.Group | measure-object Length -sum).sum}} |
   Sort Count -Descending |
   ConvertTo-HTML -Fragment -PreContent "<h2>30-90 Days</h2>"

   $180days = $files |
   where { ($_.LastWritetime -le (Get-Date).AddDays(-90)) -and
       ($_.LastWritetime -ge (Get-Date).AddDays(-180)) } |
   Group-object {if ($_.extension) {
       $_.Extension.Substring(1)}} | Select Name,Count,
```

Listing 1: *continued*

```

@{Name="Size";Expression={
    ($_.Group | measure-object Length -sum).sum}} |
Sort Count -Descending |
ConvertTo-HTML -Fragment -PreContent "<h2>90-180 Days</h2>"

$1yr = $files |
where { ($_.LastWritetime -ge (Get-Date).AddDays(-356)) } |
Group-object {if ($_.extension) {
    $_.Extension.Substring(1)}} | Select Name,Count,
@{Name="Size";Expression={
    ($_.Group | measure-object Length -sum).sum}} |
Sort Count -Descending |
ConvertTo-HTML -Fragment -PreContent "<h2>365 Days</h2>"

$summary = $groupExt | Select Name,Count,
@{Name="Size";Expression={
    ($_.Group | measure-object Length -sum).sum}} |
Sort Size -descending |
ConvertTo-HTML -Fragment -PreContent `
"<h2>Report by File Extension $Path</h2>"

# Create the HTML report.
ConvertTo-Html -head $Head `
-title "Extension Report for $Path" `
-PostContent `
($summary + $30days + $90days + $180days + $1yr) |
Out-file c:\work\extrpt.htm

```

Here are the key parts of this script:

- The code in callout A is grouping on a custom property that essentially drops the leading period from the extension name.
- The code in callout B analyzes all the files and builds new groups based on the file age. Each file age group is converted to an HTML fragment.

- The code in callout C assembles all the fragments into a single HTML report.

You can download FileExtensionAgeHTML by clicking the *Download the Code* button. This script works with PowerShell 2.0 and later. Because FileExtensionAgeHTML is a demonstration script, it has hard-coded values for the path to search and the name of the HTML file. You'll need to revise these values accordingly. I included instructions on how to do so as well as other helpful comments in the download file.

Gain Insights Without Extensive Scripting

Being able to group objects based on some criteria offers insights on an environment that might not have been possible before without extensive scripting. Group-Object doesn't care what type of object it uses. I've been demonstrating with files, but you could just as easily group Active Directory (AD) user objects or Microsoft IIS websites. Once you learn how to use Group-Object, you can apply it just about anywhere. ■

**Download**[Download the code](#)

Getting Started with System Center 2012 Orchestrator

Automate complex tasks via runbooks



John Savill

is a Windows technical specialist, an 11-time MVP, and an MCSE for Private Cloud and Server Infrastructure 2012. He's a senior contributing editor to *Windows IT Pro* and his latest book is *Mastering Hyper-V 2012 R2 with System Center and Azure* (Wiley).

Email



Twitter



Website



Blog



No conversation about Microsoft [System Center](#) 2012 would be complete without mention of the Orchestrator component, which ties other components together or integrates them with some external system. System Center 2012 Orchestrator is the new name for the product formally known as Opalis, which Microsoft acquired in late 2009.

In short, Orchestrator is an IT process-automation toolbox. It allows connectivity to practically any IT system in an organization and then performs actions (known as *activities*) on those systems. Various activities are linked together to define a complete process, which is known as a *runbook*.

Consider a manually intensive task that you currently perform—one that involves the use of multiple consoles to manage multiple systems. With Orchestrator, you can create a single runbook with activities that perform actions on all those systems; by doing so, you can automate the entire task. It gets even better when you consider that the runbooks can be called from an Orchestrator web portal; from components such as System Center 2012 Service Manager, as part of its central Service Catalog; and even from other products or command-line interfaces (CLIs), via Windows PowerShell and API mechanisms. This flexibility makes Orchestrator's capabilities readily available from other environments.

Orchestrator Architecture

Orchestrator has a simple architecture. In addition to several tools to manage the environment, Orchestrator consists of three primary components:

- **Management server**—The management server component provides a communication link between the Orchestrator database and runbook servers. The management server also deploys runbooks and integration packs to runbook servers. Only one management server can be deployed per Orchestrator deployment. If the management server is unavailable, then the tools to deploy and control runbooks (e.g., Runbook Designer) will not function. Use virtualization to make the management server resilient to server failure.
- **Runbook server**—The runbook server runs the runbooks that are deployed to it. Runbook servers communicate directly with the Orchestrator data store to determine whether a runbook needs to be run. The management server does not need to be available for runbooks to be actioned. Having multiple runbook servers is common in a highly available deployment. Each runbook server can run 50 runbooks concurrently. If more than 50 runbooks are required, then multiple runbook servers will definitely be needed.
- **Data store**—Orchestrator uses Microsoft SQL Server for its data store, which contains all deployed runbooks, configuration, log files, and system status. In a highly available configuration, the data store should be part of a clustered SQL Server deployment.

In a lab environment, the management and runbook servers can be installed on the same OS instance, such as a single virtual machine (VM). In production environments, these servers are typically split over multiple OS instances. In terms of requirements, Orchestrator has the same requirements as the rest of the System Center 2012 SP1 components, which are as follows:

- [Windows Server 2012](#) or [Windows Server 2008 R2](#)
- [SQL Server 2012](#) or [SQL Server 2008 R2](#)

Unlike the other System Center 2012 components, Orchestrator is still a 32-bit application. This isn't a problem because 64-bit OSs can run

32-bit applications. However, in some instances when calling 64-bit components, you need to do a bit of “magic” to make everything work. The good news is that because Orchestrator is 32-bit, its tools can be run on both 32-bit and 64-bit client OSs.

The core Orchestrator installation is only about 100MB. Compared with the rest of System Center, Orchestrator is tiny and has low memory and processor requirements. However, those requirements depend on the complexity of the runbooks that are executed on the runbook servers.

Orchestrator Tools

Several tools are used with Orchestrator to create runbooks, configure servers, and monitor the environment. These tools can be deployed to Windows 7, Windows 8, Windows Server 2008 R2, or Windows Server 2012. I’m going to introduce the tools as we use them to actually create and use our first runbook. I’m assuming that all the Orchestrator components and tools have been installed on a single server and that you’re logged on to that server so that they are readily available.

Creating a Runbook

I previously discussed activities, which perform actions on the systems for which Orchestrator is automating processes. Out of the box, Orchestrator has several built-in activities, which fall into the following categories:

- System
- Scheduling
- Monitoring
- File Management
- Email
- Notification
- Utilities
- Text File Management
- Runbook Control

Using only the built-in activities in these categories, you can create powerful runbooks. For example, within the System category is a Run .Net Script activity that enables PowerShell to be called, and almost anything can be accomplished with PowerShell. But the goal of Orchestrator is to make creating runbooks—and therefore automating processes—as simple and intuitive as possible. So instead of constantly creating PowerShell scripts, you'll want to take advantage of Orchestrator integration packs.

An integration pack is a collection of activities that have some commonality: They all relate to a certain product or enable a certain capability. For example, integration packs are available for each System Center 2012 component, Windows Azure, VMware, Microsoft Exchange, Active Directory (AD), HP Service Manager, and so on. Several [free integration packs are available](#) from Microsoft. Others have been created by the [Orchestrator community](#), and certain vendors supply integration packs for their solutions so that those solutions can be used in runbooks. Some companies, such as [Kelverion](#), sell these integration packs as well as having free ones available.

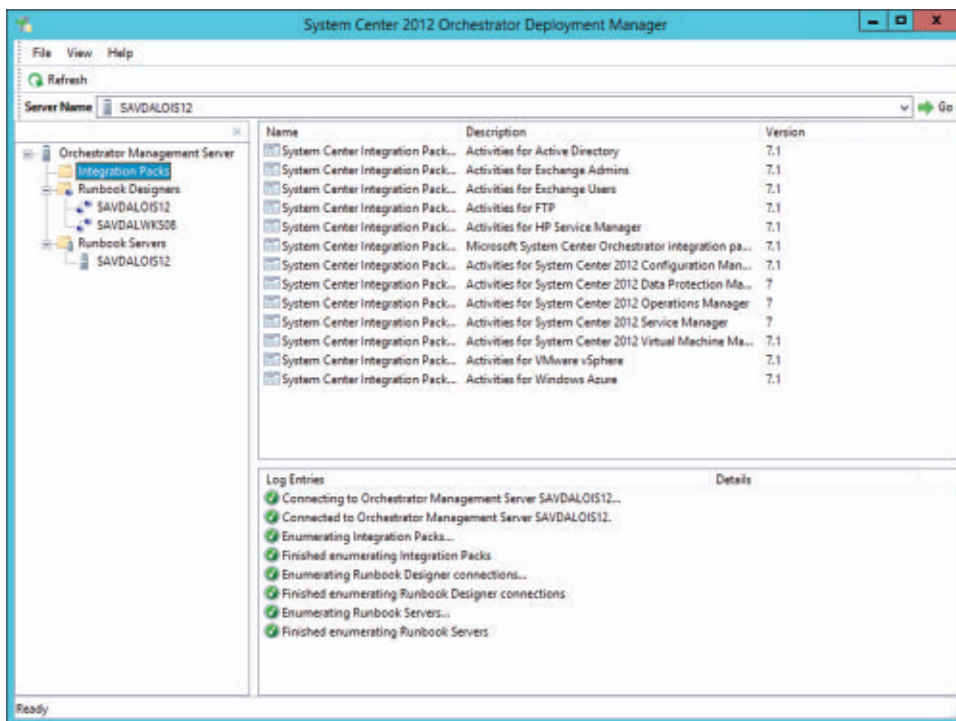
After you download an integration pack, the next step is to register it with Orchestrator. You then deploy the integration pack to Runbook Designer instances (so that the integration pack can be used as part of runbooks) and to runbook servers. The Deployment Manager tool is used to perform this function.

Launch Deployment Manager. The tool splits into three sections, showing the integration packs that are available to the Orchestrator deployment, the deployed Runbook Designer instances, and the deployed runbook servers, as Figure 1 shows.

This makes Deployment Manager not just a useful tool for deploying integration packs but also for deploying Runbook Designer instances and new Runbook servers. The tool is also a single point via which the entire Orchestrator deployment status can be viewed.

For now, we already have Runbook Designer and a runbook server on the local server. We'll use Deployment Manager to load a new

Figure 1
Deployment Manager



integration pack and deploy it to Runbook Designer and a runbook server instance. In this example, we've downloaded and will deploy the *Activities for Active Directory* integration pack.

1. Right-click the Integration Packs navigation node and select *Register IP with the Orchestrator Management Server*.
2. Click Next on the Welcome page.
3. Click Add in the integration pack selection wizard, select the OIP file, and click Open. (If you were loading more than one pack, you'd repeat these actions for each one.) Click Next.
4. Click Finish.

The integration pack is now imported into the Orchestrator deployment and registered with the management server. The next step is to deploy the integration pack to Runbook Designer and the runbook server.

1. Select the integration pack from the list displayed in the Integration Packs node, and choose the *Deploy IP to Runbook Server or Runbook Designer* action, as Figure 2 shows.

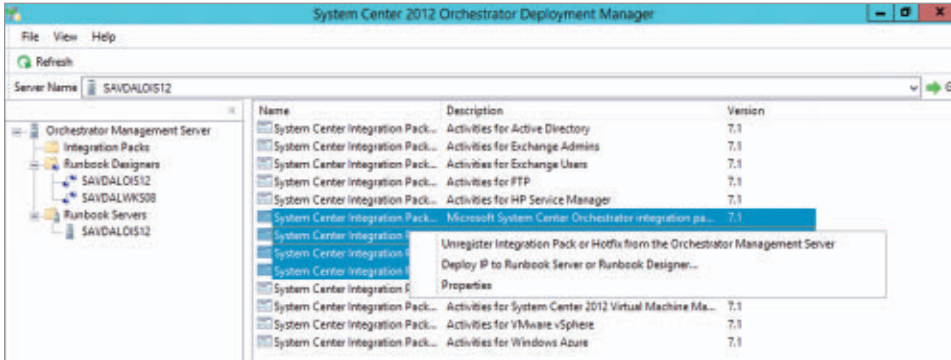


Figure 2
Deployment
Integration Packs

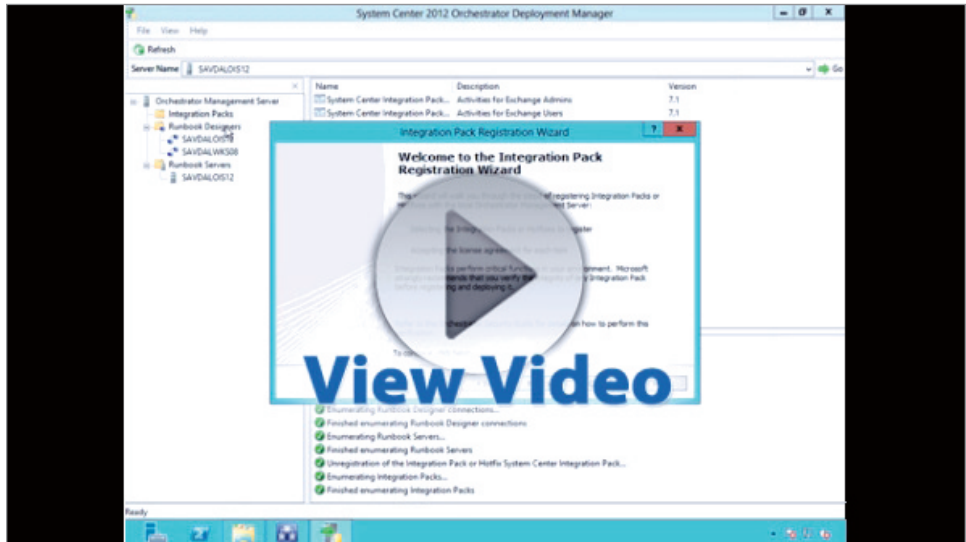
2. Click Next on the integration pack deployment wizard introduction page.
3. Choose the integration pack that you want to deploy, then click Next. (You can choose more than one pack when you're deploying more than one.)
4. Select the machines to which to deploy the integration pack (or packs); these can be Runbook Designer systems or runbook servers. Click Next.
5. Choose whether to perform the deployment immediately, or schedule the deployment for a set time. Click Next.
6. Review the list of actions that will be performed, then click Finish.

The integration pack is now available for inclusion in runbooks and for use on runbook servers. As you navigate through Deployment Manager, note that if you right-click the Runbook Designers or Runbook Servers navigation node, the option to deploy a new instance is available. This option makes it easy to expand your Orchestrator deployment. The whole process can be seen in the accompanying video.

Video



John Savill demonstrates the Active Directory Integration Pack for System Center Orchestrator 2012 SP1



Now it's time to use Runbook Designer to create and test runbooks. Launch Runbook Designer. Note the navigation on the left side, which shows the runbooks that exist, as well as the options to create computer groups and deploy runbooks to runbook servers. The right side of the tool shows all the integration packs and built-in activities that are available. The majority of the tool comprises the canvas on which you drag activities to create runbooks.

The first step is to create a new folder for your runbooks. We'll then create a new runbook within that folder.

1. Right-click the Runbooks navigation node and choose New, Folder.
2. Enter a name (e.g., ITPRO), then press Enter.
3. Right-click the new folder and choose New, Runbook.
4. A New Runbook tab is created. Right-click this tab, choose Rename, and give the tab a meaningful name (e.g., NewUser). To complete the renaming process, the runbook needs to be checked out. This checkout is automatically performed for you.

In this example, we're creating a runbook that we'll use to create user accounts. To do this, we'll use the Active Directory integration pack.

Some integration packs use configurations that need to be created before the activities can be used. For the Active Directory pack, you must create a configuration for the AD instance with which you want to interact. This step requires the name of a domain controller (DC), a credential, and the default container in which objects will be created.

1. Under the Options menu in Runbook Designer, you'll see all the integration packs that require created configurations. Choose Active Directory.
2. Click the Add button.
3. Fill in the available fields. The parent container needs to be in distinguished name (DN) format. For example, the default Users container for my savilltech.net domain would be CN = Users,DC = Savilltech,DC = Net. When all the fields are completed, click OK.

You can now create the runbook.

1. The runbook should already be checked out from the renaming process. If it isn't, choose the Check Out action on the main menu bar.
2. In the Activities area, select Runbook Control, then drag the Initialize Data activity onto the canvas, on the far left. We'll use this activity to define the variables that will be used throughout the runbook.
3. Double-click the Initialize Data activity on your canvas; Details is selected.
4. Click Add to create new parameters, then click the placeholder name of each created parameter and give it a unique name. In our example, we want to add three parameters; name them CommonName, FirstName, and Surname, then click OK. These values are required when running the runbook and will be placed on the Orchestrator data bus, which is available to all the activities in the runbook and is a great way to pass data.

The next step is to add an activity to create a new user.

1. In the Activities list, select Active Directory, then drag the Create User activity onto the canvas, to the right of the Initialize Data activity.
2. Hover over the canvas, just to the right of the Initialize Data activity. A small arrow will appear. Click and hold the arrow, then drag it over the Create User activity. Doing so creates a link between the two activities.
3. Double-click the Create User activity, then click the button to the right of the configuration name and select the configuration that you created for the Active Directory integration pack.
4. Under Properties, the CommonName parameter is shown. Right-click the blank space next to the CommonName value. We will use the value from the Initialize Data activity by subscribing to its data on the Orchestrator data bus.
5. From the context menu, choose Subscribe, Published Data. The Initialize Data activity is selected by default. Choose the CommonName value that's shown and click OK.
6. In this example, we also want to set the first name and surname of the new user. Click the Optional Properties button, choose Last Name and First Name, and click the right-arrow button to add to the properties to be used.
7. Use the Subscribe action again for each value, to select the corresponding value from the Initialize Data activity in the same way you did for CommonName in step 5. Click Finish.
8. These actions create a disabled user. To enable the user, we need to add the Enable User activity. Drag the Enable User activity from the Active Directory integration pack onto the canvas, to the right of the Create User activity. Create a link from the Create User activity to the Enable User activity, as described in step 2. If you right-click the link, you'll see the default Include option, which means that the link is used if the Create

User activity returns success. You could link other activities that can be used if Create User fails.

9. Double-click the Enable User activity and choose the AD configuration. Right-click anywhere in the value space and select Subscribe, Published Data. For Activity, choose Create User, then choose the Distinguished Name value (which is populated automatically as part of user creation).
10. Click Finish. The finished runbook should look like the one that Figure 3 shows.

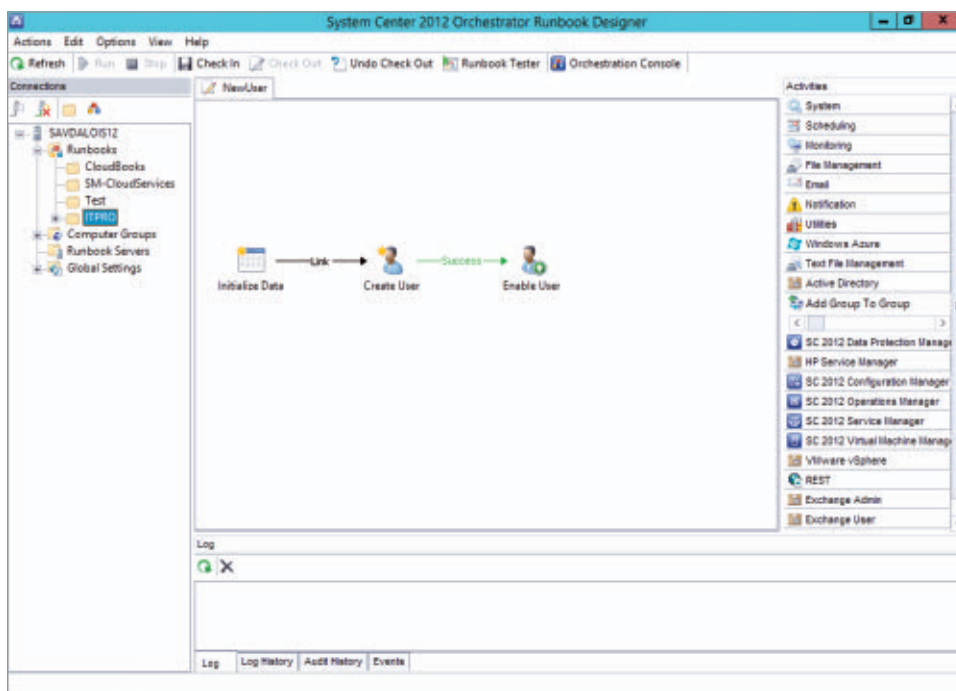


Figure 3
Creating a
User Runbook

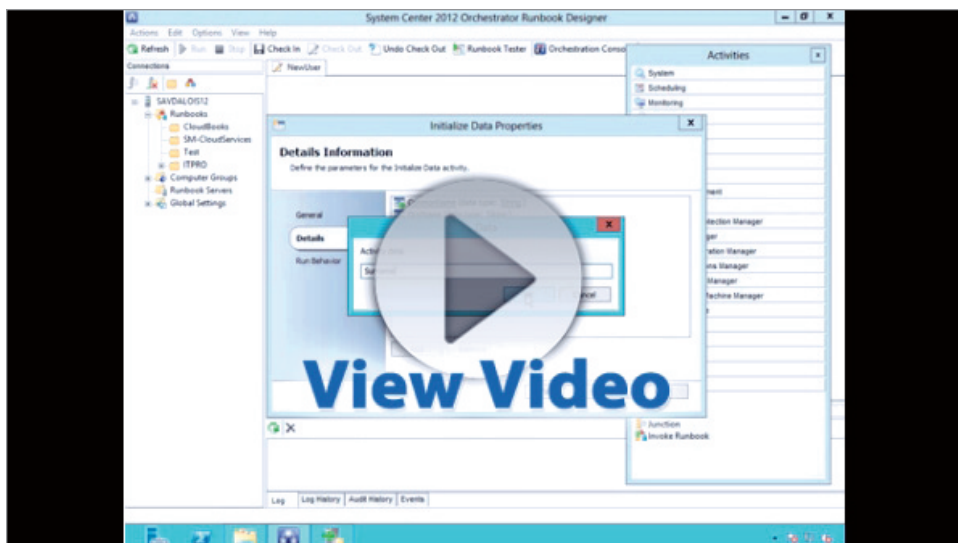
The created runbook can now be used. To test the new runbook, click the Runbook Tester button, which opens the Runbook Tester utility. Click the Run button. You'll be prompted for the values that are required as part of the Initialize Data activity. After you complete these values, the runbook will run and you can see the status of each activity in it.

In this example, I used only activities from the Runbook Control and Active Directory integration packs. The real power comes from using activities from many integration packs, allowing the automation of processes that require the use of many systems. Various levels of logging can be enabled to help track runbook usage and troubleshoot possible problems. This entire process is shown in the accompanying video.

Video



John Savill explains how to create a new runbook in System Center 2012 Orchestrator



Now that the runbook has been created, it's available by default for runbook servers to utilize. The new runbook is also available in the Properties of each runbook server. You can specify runbook servers that will be used to execute the runbook. Although runbooks can be executed through many different means, Orchestrator includes various web options, such as a complete Microsoft Silverlight web-based console, which by default is available on port 82 of your management server. (The console can also be launched from Runbook Designer, through the Orchestration Console button.) The console allows full tracking or initiation of runbooks, as shown in Figure 4.

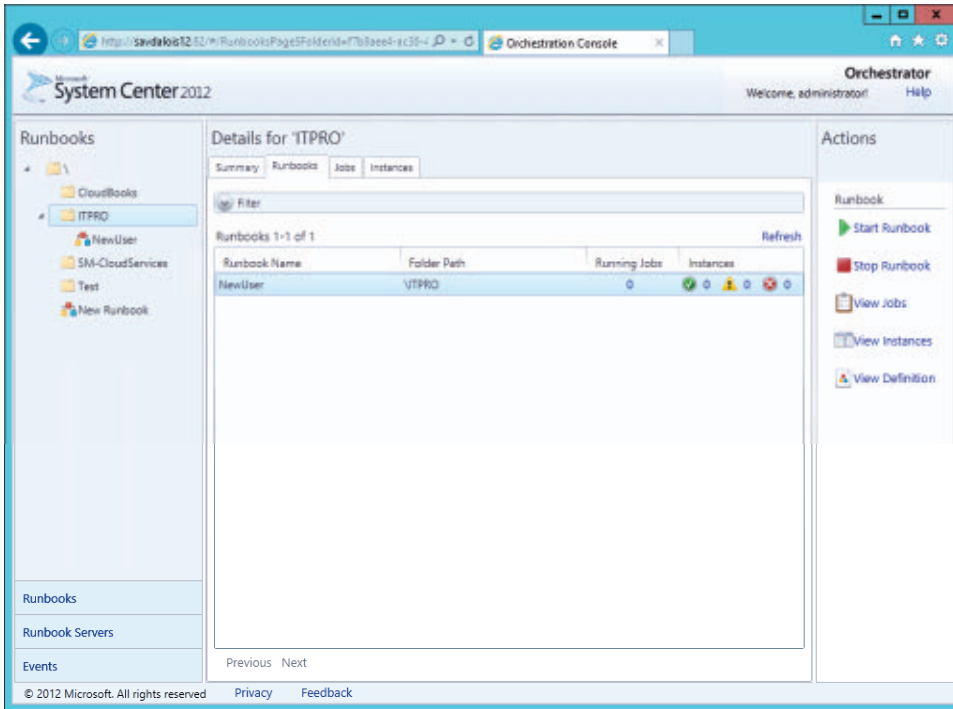


Figure 4
Silverlight Web-Based
Console

Start Small, Gain Big

The best way to get started with Orchestrator is to start small. Think of a process that involves some manual effort, and create a runbook to automate that task. You might even want to convert a PowerShell script into an Orchestrator runbook. When you start to use Orchestrator, you'll quickly learn how easy it is to automate what could otherwise be a very complex process. When you have runbooks, integrate them into other components, such as Service Manager, to make them easily available. (I touched on this topic in the article [“What’s New with System Center 2012 Service Manager SP1.”](#))

One final tidbit: Download the [Best Practices Analyzer](#) and use it in your environment. This tool can quickly point out any possible problems and help your environment run as efficiently as possible. ■

FAQ

Answers to Your Questions



John Savill

Q: Since I have multiple virtual networks in Windows Azure, can I make a virtual machine accessible to all virtual networks?

A: Each virtual network in Windows Azure can have a single site-to-site VPN configured that can link back to your on-premises environment. If you had multiple virtual networks in Windows Azure and a virtual machine (VM) in one of those virtual networks that you wanted to make accessible to all the virtual networks, you would need to have a site-to-site VPN from each virtual network link back to your on-premises network, then route the traffic between the virtual networks using your on-premises network. At this time, there's no way to route locally in Windows Azure between virtual networks, and a VM can be part of only one virtual network.

—John Savill



Jan De Clercq

Q: What happens to a user account's password when I select the *Smart Card is required for interactive logon* option in the user's Active Directory account properties?

A: When you select the *Smart Card is required for interactive logon* check box in the Active Directory (AD) user account properties, Windows automatically resets the user password to a random complex password. In addition, Windows adds the SMARTCARD_REQUIRED flag to the UserAccountControl user account attribute and sets the DONT_EXPIRE_PASSWORD flag on the user account. The

latter ensures that the user's password never expires after the *Smart Card is required for interactive logon* option is selected.

When a user logs on to Windows either locally or remotely using a Remote Desktop session, the Windows client automatically checks for the presence of the SMARTCARD_REQUIRED flag. If the *Smart Card is required for interactive logon* option is set for the user, Windows rejects the logon attempt if it's not made with smart card credentials.

—Jan De Clercq

Q: With Hyper-V Replica, how can I store different replicas in different locations on the target?

A: When configuring Hyper-V Replica, a location is configured on the target server, but it's not possible to specify different locations on a per virtual machine (VM) level. The solution is to enable replication. After the replication is established, use storage migration to move the replicated VM to another storage location.

—John Savill

Q: I am trying to connect to my System Center 2012 R2 Orchestrator Web Console or the Web Service, but I receive errors—what can I do?

A: The most common cause for this problem is that the owner on the databases used for Orchestrator services isn't Orchestrator. This causes various types of errors for the Orchestrator Web Console and Web Service. The easiest way to solve this problem is to do the following:

1. Log on to SQL Server.
2. Launch SQL Server Management Studio (SSMS).
3. Expand Databases, Orchestrator, Security, Users.
4. Right-click the Orchestrator service account and select Properties.

5. Select the Membership tab on the displayed dialog box.
6. Select the db_owner role membership, and click OK.

The Orchestrator web components should now work.

—John Savill

Q: How can I find out if a smart card was used to log on to Windows? Are there specific Windows event-log entries I can scan for?

A: The Security Event log in a Windows domain controller (DC) provides entries that you can use to detect smart card logons. In the log, you must scan for successful Account Logon events that have the ID 672. These entries signal a successful Kerberos authentication ticket grant. Event 672 records who requested the Kerberos ticket, the client's IP address, and the type of authentication credentials in the Pre-Authentication Type field. When a smart card was used, the Pre-Authentication Type field shows the value of 14, 15, 16, or 17. Under the hood, these values refer to PKINIT protocol messages. PKINIT is the Kerberos protocol extension that Windows uses for enabling smart card logons. It stands for “**P**ublic **K**ey Cryptography for **I**nitial Authentication.” For the detailed syntax of event 672, see the [TechNet support page for this event](#).

—Jan De Clercq

Q: How can I quickly find all the various port requirements for different Microsoft products?

A: Microsoft maintains a complete list of all the ports used for its major products. You can see it at the Microsoft site “[Service overview and network port requirements for Windows](#).” Find and select the product, and it will show all the various port needs, which can then be used to help in firewall configuration or other

tasks. System Center Virtual Machine Manager (SCVMM) has its own specific port requirements shown at “[VMM Ports and Protocols](#).”

—John Savill

Q: What can we do to limit or exclude the use of the RC4 stream cipher on our Windows platforms? What are the Microsoft recommendations for disabling RC4?

A: Microsoft recommends that customers use Transport Layer Security (TLS) 1.2 and the more secure Advanced Encryption Standard - Galois/Counter Mode (AES-GCM) cipher as the RC4 alternative. You can find more information about this recommendation in the TechNet blog “[Security Advisory 2868725: Recommendation to disable RC4](#).”

Internet Explorer 11 (IE 11), which is bundled with [Windows 8.1](#), enables TLS 1.2 by default and no longer uses RC4 during the SSL/TLS handshake. More details about this can be found in the MSDN blog “[IE11 Automatically Makes Over 40% of the Web More Secure While Making Sure Sites Continue to Work](#).”

Microsoft also released a patch that provides support for the IE 11 and Windows 8.1 RC4 changes on Windows 8, [Windows 7](#), Windows RT, [Windows Server 2012](#), and [Windows Server 2008 R2](#). You can find more information about the patch in the Microsoft Support article “[Microsoft security advisory: Update for disabling RC4](#).”

The RC4 cipher can be completely disabled on Windows platforms by setting the “Enabled” (REG_DWORD) entry to value 00000000 in the following registry locations:

- HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\SecurityProviders\SCHANNEL\Ciphers\RC4 128/128
- HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\SecurityProviders\SCHANNEL\Ciphers\RC4 40/128
- HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\SecurityProviders\SCHANNEL\Ciphers\RC4 56/128

Windows clients that have these registry entries set won't be able to connect to sites that require RC4. Windows servers that have these registry entries set won't be able to service clients that must use RC4.

—Jan De Clercq

Q: How do I enable processor compatibility on Windows Server 2012 Hyper-V for earlier OSs?

A: Early versions of Hyper-V offered the option for processors to support older, legacy OSs such as Windows NT 4.0. This configuration hid additional processor information from the virtual machine (VM) such as logical processor to cache mapping, which otherwise would cause the earlier OSs to crash. The ability to hide this is still available and is configured by using Windows PowerShell:

```
Set-VMProcessor -CompatibilityForOlderOperatingSystemsEnabled
    $true
```

—John Savill

Q: Can I use a Windows Azure virtual machine with 16 1TB data disks attached as a single volume?

A: Windows can create striped volumes that effectively join multiple disks together. Add the disks to the Azure virtual machine (VM), then log on to the VM and start the disk management tool, diskmgmt.msc. Select one of the disks and pick the Create Striped Volume action, ensuring all disks are included. Choose the quickformat option to minimize format time. When it's done, you'll have a single 16TB volume to use. For Linux OSs, use the MD capability or LVM to get the stripe. ■

—John Savill

Product News for IT Pros

iPass Launches Windows Phone 8 Mobile App for Wi-Fi Connectivity

iPass announced iPass Open Mobile for [Windows Phone 8](#), which enables Wi-Fi connectivity for users of Windows Phone 8, alongside its existing Android and iOS capabilities. The iPass Open Mobile solution is ideal for keeping mobility costs down and mobile worker productivity up in the Bring Your Own Device (BYOD) era. Available as a free download from the Windows Store, iPass Open Mobile for Windows Phone 8 provides employees with seamless connectivity to the world's largest commercial Wi-Fi network of more than 1.4 million hotspots, including in-flight Internet, hotels, airports, and business venues around the world. iPass Open Mobile on the Windows Phone 8 platform lets you provide employees with seamless connectivity to the iPass Mobile Network, providing a consistent user experience while avoiding costly roaming charges and Wi-Fi session passes. Organizations gain visibility into enterprise, group, individual, and device mobility and can easily enforce expense and security policies—helping to reduce the total cost of mobility. For more information, visit the [iPass website](#).



iWeb Launches Microsoft Private Cloud Hosting

iWeb announced its new [Microsoft Private Cloud](#) service, designed to give businesses complete control and flexibility of their IT infrastructure. Businesses can now choose the Microsoft Private Cloud to quickly add or scale up virtual servers as needed, without the expense or complexity of deploying their own data centers and network infrastructure. The Microsoft Private Cloud from iWeb lets businesses use virtualization to build a software-defined data center hosted at iWeb



but managed by their own IT teams. iWeb's dedicated infrastructure is fully managed in the company's state-of-the-art data centers and delivered on one of the world's most reliable networks, with guaranteed service level agreements (SLAs) and 24 × 7 support. With the Microsoft Private Cloud from iWeb, businesses can extend their in-house data center or move specific workloads to a hosted environment; alleviate concerns about in-house data center space, power, and cooling; eliminate capital expenditure for storage infrastructure; and take advantage of on-demand flexibility with virtualization and cloud technologies. For more information, visit the [iWeb website](#).



Riverbed Granite 2.6 Provides Instant Branch Recovery

Riverbed Technology announced Riverbed Granite 2.6, the latest version of the company's branch-converged infrastructure solution, which centralizes branch data in the data center while delivering local performance to branch users. With Granite, businesses can restore operations in a matter of minutes instead of days, centrally protect and secure data, and significantly lower the TCO of branch and remote offices. New features include enhanced snapshot support for enterprise-class storage solutions, application-consistent data protection with a greater number of data center storage arrays; integrated IBM Storwize V7000 snapshot support; a snapshot handoff framework that introduces a script-execution interface used to orchestrate snapshot operations for storage arrays; and higher-capacity Virtual Granite Core (VGC) models that delivers greater capacity for VGC deployments, scaling to support more branches and larger datasets. Learn more at the [Riverbed Technology website](#).



AvePoint DocAve Online SP3 Offers New Data Protection, Governance, and Reporting

AvePoint announced the availability of DocAve Online SP3, an update that features data protection, governance, and reporting enhancements that empower organizations to maintain the same level of

protection and control over their cloud-based assets as they have with on-premises solutions. Updates in DocAve Online SP3, AvePoint's SaaS platform for Microsoft Office 365 management, include policy enforcement, ensuring that all sites remain within established IT governance policies by monitoring sites around the clock; granular content protection, enhancing business continuity by restoring SharePoint Online, SkyDrive Pro, and Exchange Online content to any supported storage locations; and configuration reports that give you deeper insight into the graphical topology of SharePoint Online for streamlined architecture planning. For more information, visit the [AvePoint website](#).

RoboForm Password Manager for Windows Phone



Siber Systems announced the availability of RoboForm for Windows Phone, a free app that extends the power of RoboForm to smartphones running Windows Phone. RoboForm is a powerful password manager that lets users securely log on to banking, ecommerce, and healthcare sites as easily as using a browser bookmark. In addition to eliminating the need to remember dozens of usernames and passwords, RoboForm stores identities, bookmarks, and contacts. Users can auto-generate different usernames/passwords for each use; the information is then stored in a RoboForm Everywhere account to be used from the RoboForm for Windows Phone app. For more information, visit the [RoboForm website](#).

Barracuda Introduces Powerful F280 Desktop Appliance



Barracuda Networks announced the availability of Barracuda NG Firewall F280, which brings next-generation firewall capabilities and full gigabit Ethernet throughput to branch offices and remote locations. The Barracuda NG Firewall is an enterprise-grade network firewall that offers massive scalability, easy and efficient management across dispersed networks, low resource consumption, and high performance for business-critical applications. The Barracuda NG Firewall F280

provides all the features of the Barracuda NG Firewall version 5.4.2 release, including application detection, SSL interception, and malware protection; direct, fast, and secure access to Internet and cloud resources; 99.99 percent availability of all business-critical resources; integrated application-aware WAN virtualization; integrated wireless for mobile devices; and central management using Barracuda Control Center. The Barracuda NG Firewall is available as a hardware and virtual appliance. To learn more, visit the [Barracuda Networks website](#).



Spiceworks Debuts New Security Integrations and Help Desk Plug-ins

Spiceworks announced Spiceworks 7.1, an update to the company's September 2013 release of Spiceworks 7. The update includes performance enhancements to its IT inventory, Help desk, and mobile device management capabilities, as well as new ways for technology vendors to more seamlessly integrate their offerings with Spiceworks. The company also debuted four new Help desk plug-ins that give IT pros the ability to customize their Spiceworks Help desk system to suit their needs. New security integrations include AlienVault Threat Alerts, which help users identify and mitigate security threats on their network for free; NetClarity, which lets Spiceworks users block and/or take action against unwanted devices attempting to access their networks; and Webroot, an integration of Webroot's SecureAnywhere that enables IT pros to receive real-time threat alerts and view and manage endpoint security directly from Spiceworks. For more information, visit the [Spiceworks website](#). ■

Search our network of sites dedicated to hands-on technical information for IT professionals.
www.windowssitpro.com

Support

Join our discussion forums. Post your questions and get advice from authors, vendors, and other IT professionals.
forums.windowssitpro.com

News

Check out the current news and information about Microsoft Windows technologies.
www.winsupersite.com

EMAIL NEWSLETTERS

Get free news, commentary, and tips delivered automatically to your desktop.

- Cloud & Virtualization UPDATE
- Dev Pro UPDATE
- Exchange & Outlook UPDATE
- Security UPDATE
- SharePoint Pro UPDATE
- SQL Server Pro UPDATE
- Windows IT Pro UPDATE
- WinInfo Daily UPDATE

RELATED PRODUCTS

Windows IT Pro VIP

Get exclusive access to over 40,000 articles and solutions on CD and via the web. Includes FREE access to eBooks and archived eLearning events plus a subscription to either *Windows IT Pro* or *SQL Server Pro*.
windowssitpro.com/vip-premium-membership

SQL Server Pro

Explore the hottest new features of SQL Server, and discover practical tips and tools.
www.sqlmag.com

Dev Pro

Discover up-to-the-minute expert insights, information on development for IT optimization, and solutions-focused articles at DevProConnections.com, where IT pros creatively and proactively drive business value through technology.
www.devproconnections.com

SharePoint Pro

Dive into Microsoft SharePoint content offered in specialized articles, member forums, expert tips, and web seminars mentored by a community of peers and professionals.
www.sharepointpromag.com

Advertiser Directory

Kroll Ontrack.....	2
ManageEngine	6
Netwrix	7
PASS.....	8
StorageCraft	1
Windows IT Pro.....	43

Vendor Directory

Adobe	19, 20
AvePoint.....	88, 89
Barracuda Networks	89, 90
B-side Software	19
Evernote.....	20

Facebook.....	18, 19
Flipboard	19
HP.....	73
iPass.....	87
iWeb.....	87, 88
Kelverion	73
Next Matters	19
Nokia	14
Riverbed Technology.....	88
Siber Systems	89
Spiceworks.....	90
Twitter	19
VMware	73
Wikipedia.....	20